Information and the disposition of medical malpractice claims: A competing risks analysis*

Paul Fenn
University of Nottingham

Neil Rickman
University of Surrey
Erasmus University
and CEPR

December 21, 2011

Abstract

We explore empirically, using a competing risk model, the relationship between information about case strength and the speed with which medical malpractice disputes are resolved. We have data on the time to resolution of a number of such disputes in a group of English hospitals, as well as the means by which each dispute is resolved (drop, settlement, or trial). In addition we have detailed data on the evolution of expert assessments of case strength, and on the timing of procedural events (i.e. external experts' reports) that are designed to share information and that, therefore, might be expected to influence litigation outcomes. We find that litigation encourages dropping and settling of cases over time in a systematic way relating to their assessed strength; cases that involve relatively little uncertainty are resolved faster than those where liability appears to be more unclear. We suggest that this evidence is consistent with the litigation process using time to help sort, and deal with, cases according to their strength.

JEL number: C7, K4

Keywords: Information; Delay; Competing risks; Medical malpractice.

*We are grateful to the Editor and two anonymous referees for comments. We are also grateful to Parimal Bag, Winand Emons, Jennifer Reinganum, Yasutora Watanabe, Micelle White and participants at the European Association of Law and Economics Annual Conference, Zagreb, September 2003; the American Law and Economics Annual Conference, New York, May 2005; the Legal Services Workshop at the CMPO, Bristol, March 2007; and the Law and Economics Workshop, Taormina, March, 2007 for helpful comments on various drafts of the paper. Remaining errors are our own.

1 Introduction

What determines the course of a legal claim as time elapses? At any one point in time, there are several possible decisions faced by the litigants—the plaintiff could decide to drop the claim; the parties could agree a settlement, or they could decide to continue the dispute into the 'next period'. The timing of the claim's resolution, and the nature of that resolution, are therefore the outcomes of choices made between competing alternatives as long as the claim remains live. In this paper we are interested in the factors that dictate these choices: are there systematic reasons for some cases to be dropped early, or to settle late (or vice versa)? We argue that this is important for two reasons. First, the duration of claims (or the 'delay' to resolution) is a long-standing policy concern across many jurisdictions.[1] Delay is often accompanied by high legal expenditures, uncertainty for the parties and intertemporal redistribution from plaintiff to defendant: each of these may affect incentives to bring claims and, hence, may inhibit the efficiency (deterrence) and equity (compensation) objectives of the legal system. Second, despite these costs, a potential benefit of delay is that it may help to sort weak from strong claims, by helping the parties to interpret evidence and produce new information. This may encourage mutually acceptable outcomes and help to filter relatively clear-cut cases away from a costly trial.[2] It is of fundamental concern, therefore, to discover the extent to which both drop and settlement decisions contribute to the sorting process over time.

From a theoretical perspective, there are a number of different approaches to modelling the reasons why a case may settle, drop or go to trial: the parties

[1]See, for example, the evidence cited in Fenn and Rickman (1999).

[2]Farber and White (1991)'s study of hospital informal dispute resolution suggests such a filtering mechanism, as does Priest and Klein (1984)'s analysis of cases that reach trial.

may view common information differently ('divergent expectations models') or they may possess private information and behave strategically ('asymmetric information models').[3] In turn, as we indicate in Section 2, these models make different predictions about the types of cases that settle or go to trial, while they generally fail to take simultaneous account of the option to drop the case during the litigation, as well as settle or continue it. A number of empirical studies have also examined drop and settlement behaviour independently, broadly following the trajectory of theoretical work on litigation. Thus, Danzon and Lillard (1983), Viscusi (1986, 1988), Sloan and Hoerger (1991) and Farber and White (1991, 1994) examine the factors that affect the probability of a claim being dropped in a static (probit) estimation framework. A number of these studies, plus instances like Fournier and Zuehlke (1989) and Sieg (2000), estimate factors determining the probability of settlement. In a dynamic setting, Fournier and Zuehlke (1996), Kessler (1996) and Fenn and Rickman (1999, 2001) all estimate settlement hazard functions to examine settlement behaviour over the course of a claim.

The starting point for our paper is that the decision to drop and the decision to settle are essentially competing risks (choosing one precludes the subsequent choice of the other) and should be treated as such in estimation. Doing this provides a richer insight into the factors leading to claim resolution over time. Several empirical studies have recognised the 'jointness' of settlement and drop decisions: Danzon and Lillard (1983) and Fournier and Zuehlke (1989) both estimate structural static models to capture this while

[3]We can also distinguish different approaches to the modelling of time in litigation, with divergent expectations models being essentially timeless (see Landes (1971), Gould (1973), Posner (1973) and Shavell (1982)), and asymmetric information ones being divided between static one-shot (Bebchuk (1984), Reinganum and Wilde (1986)) and dynamic ones (Spier (1992)). See Spier (2007) and Daughety and Reinganum (2008) for thorough surveys.

Farber and White (1991) estimate 'drop' and 'settlement amount' equations jointly (though find that they cannot identify all the parameters). None, however, have studied the relationship between these decisions over time.[4] We are able to do this through use of a unique set of data on the time to resolution of a number of medical malpractice disputes in a group of English hospitals, as well as the means by which each dispute is resolved (i.e. by the patient unilaterally withdrawing her claim, or by the hospital paying an agreed financial settlement). In addition, the dataset contains information on the severity of the injury, estimated damages and the source of plaintiff finance. We are therefore able to estimate 'cause-specific' drop and settlement hazards in a competing risks regression framework.

In the process of doing this, we consider a natural question about the process determining the settlement and dropping of claims: does it reflect the strength of the claims? As stated above, some might argue that the ability of litigation to produce information and help sort 'strong' from 'weak' claims is central to its justification. Conceptually, the availability of new information in litigation may have mixed effects. For example, Cooter and Rubinfeld (1994) and Watanabe (2006) both suggest that information improves settlement prospects, while Loewenstein and Moore (2004) present experimental evidence that the precise effects of information may depend on its dimensionality. Empirically, Huang (2009) finds evidence that the production of information (via discovery) encourages settlement. In fact, our interest in the sorting effect of litigation recognises that information may do more than affect settlement probabilities: it may also help to match outcomes to case strength—in this respect, our paper develops earlier work by Farber and

[4]Helland and Tabarrok (2003) use difference-in-difference estimators to consider patterns of drop behaviour and claim durations across US States over time; but their State-level data prevent them from tying these together for individual claims.

White (1991). Our data allow us to consider this issue because they contain contemporaneous expert estimates of case strength as it varies across claims and over time. In addition, they also record the timing of important events during the litigation process such as the arrival of new information in the form of expert reports. It is rare to have access to such detail. Probably the closest analogy to our data on assessed case strength is Farber and White (1991)'s. They have expert assessments of case strength provided by hospital management and they mix this with external expert reports. Our data improve on this by including changes made to the assessments over time and, thanks to a considerably larger sample size, by allowing us to treat external expert assessments as separate sources of information. In the absence of such data, other approaches to measuring case strength have involved the creation of proxies via *ex post* third party case file analysis (Sloan and Hoerger (1991), Studdert *et al.* (2006)) and treatment as an unobserved structural parameter to be estimated via GMM (Sieg (2000)). It is notable that previous duration analyses of litigation have lacked assessments of case strength (Fournier and Zuehlke (1996), Kessler (1996)) or changes in them over time (Fenn and Rickman (1999, 2001)).

Unlike previous papers, therefore, we are able to explore empirically, using a competing risk model with time-varying covariates, the extent to which litigation produces, and makes use of, information in determining the ultimate disposition of a claim. We believe that this is the first time data on the duration of negotiations and contemporaneous observations on information about case strength have been jointly analysed to establish how they interact with delay and generate case outcomes.

We find that, after a good deal of initial variation in the hospital's assessment of its liability, cases are ultimately resolved in a systematic way.

Thus, cases where the hospital assesses its defence to be poor settle early, while those where it judges its case to be stronger typically end with being dropped by the plaintiff. Whilst sorting these cases reasonably quickly, the litigation process uses time to sort initially less clear cases according to strength. It would seem that negotiations reveal, or credibly reinforce, views about case strength as time elapses and we find that an additional influence here is the production of new information through expert reports. We suggest that our results are consistent with a dynamic interpretation of Priest and Klein (1984)'s selection hypothesis where 'drops' are added to the plaintiff's decision space. Overall, the competing risks of drops and settlements act as twin 'blades' that remove weak and strong cases from the litigation pool. As mentioned above, we are able demonstrate the combined influence of these two effects by estimating the overall hazard of settlement. By suggesting that litigation is effective at sorting cases according to strength, our results may provide some counterweight to criticisms it has received for its costliness in many jurisdictions.

The paper is structured as follows. The next section presents a brief survey of theoretical models of litigation outcomes in order to demonstrate the variety of hypotheses about the types of case that drop, settle or go to trial and the value of empirical analysis to help tie these down. The third and fourth sections present our data and estimating methodology respectively. Section 5 presents results and remarks on some of their implications for theoretical models. Section 6 presents conclusions.

2 Theory

The potential outcome of litigated cases has received considerable attention in the literature. Whilst there is broad agreement on the effects of key variables like the cost of litigation, the role played by the strength of the case is less clear cut, with a good deal depending on the extensive form chosen to model the litigation game and the precise location of information about case strength within this. Since our empirical model examines the role of this information in litigation, we begin by demonstrating some of the possibilities on offer.[5]

Early models of litigation (e.g. Posner (1973), Shavell (1982)) argued that a case will go to trial if there is no surplus available to the litigants from settling. Thus, if the damages in question are x, the plaintiff thinks she will win at trial with probability π_p, the defendant reaches his own judgement of this (π_d) and the cost of trial to each party is c, then under UK cost rules[6] pre-trial settlement will occur when $\pi_p x - (1 - \pi_p)2c < \pi_d(x + 2c)$, i.e. when

$$\pi_p < \pi_d + \frac{2c}{x + 2c} \tag{1}$$

This says that the plaintiff must be sufficiently pessimistic about trial relative to the defendant. In this model parties are assumed to form their assessments of the case from commonly observed information that they interpret differently—they have divergent expectations. In an attempt to explain the source of these divergences, Priest and Klein (1984) argue that they are the

[5]Spier (2007) and Daughety and Reinganum (2008) provide extensive surveys of the theoretical literature on litigation.

[6]These make the losing litigant liable for the winner's costs, as well as their own. This assumption is consistent with our UK data and does not significantly alter the ensuing discussion.

result of errors made by the parties when evaluating the case.[7] They show that trial cases are a biased selection of those initially filed because, on average, they will be evenly contested and have plaintiff win rates close to 50%. Similarly, settled cases will reflect situations where the parties believe the plaintiff's case is relatively weak or relatively strong.

An alternative approach to modelling litigation is offered by Bebchuk (1984). He argues that the basis for divergences in expectations must lie in asymmetries of information between the parties and that, as a result, the bargaining extensive form should take explicit account of this. Bebchuk assumes that an uninformed party makes a take-it-or-leave-it offer which is selected to screen the informed party, with only those who are least liable going to trial. This represents an alternative prediction to that in Priest and Klein because it suggests that the more 'extreme' end of the case types will end up at trial while the most liable defendants accept the settlement offer.[8]

Spier (1992) extends the analysis to a finite number of pre-trial bargaining periods and shows that settlement will typically have a positive probability in each of these: i.e. she provides the first theoretical basis for considering the litigation process as an intertemporal one. The plaintiff's sequence of offers again means that only the strongest defendants end up at trial.

A natural concern with Bebchuk and Spier's models is the apparently arbitrary nature of the one-sided information asymmetry: what happens if both parties have private information? Daughety and Reinganum (1994) provide an early contribution but Friedman and Wittman (2007)'s approach is especially helpful for us because it considers the role of case strength in

[7]Hylton and Lin (2009) provide a good survey of work based on Priest and Klein's original paper. Farmer and Pecorino (2002) also provide a discussion of how divergent expectations may emerge.

[8]Reinganum and Wilde (1986) produce a similar type of result in a model where the informed party makes a single signalling offer.

litigation outcomes.

The authors consider a situation where litigation starts with plaintiff and defendant receiving independent private signals (π_p and π_d) about their prospects at trial.[9] They then make single simultaneous offers in an attempt to settle the case (see Chatterjee and Samuelson (1983)). Trial can arise because the two sides may be unable to fashion compatible offers given their signals. In contrast to Priest and Klein, it is now possible for cases with highly divergent beliefs to go to trial, depending on the value of c. In particular, when costs are low only a small difference between π_p and π_d is necessary to reverse (1) and yield trial; but 'small' differences are consistent with π_p and π_d both being large or small and, thus, with both agreeing that one side has a strong case.

We have so far ignored a potentially important part of the dispute resolution process: the role of cases that are dropped in helping to select those for settlement and trial. A key reason for this omission is the relatively few papers that have dealt with this interaction—not least because of the larger strategy space that it requires to be modelled.[10]

Shavell (1982) sets out the basic condition for a case to be dropped in a divergent expectations framework.[11] For this the case must have a 'negative expected value' (NEV) to the plaintiff: i.e. $\pi_p x - (1 - \pi_p)2c < 0$. Clearly, cases with high costs or weak plaintiffs are more likely to be dropped, and

[9] Friedman and Wittman model this as a signal about the value of the case, not its strength, but our version is more in keeping with the current paper.

[10] Friedman and Wittman (2007) are open about this: "Incorporating the possibility of a negative expected return [for the plaintiff] from going to trial in a two-sided incomplete information model is an interesting challenge for the future." (p. 113). Since our data provide the rare prospect of assessing the role of 'drops' it helps to assess the need for this future research, as well as the merits of existing models that ignore such strategies.

[11] In fact, since this framework is essentially timeless, Shavell describes the conditions for a case not to be filed (as opposed to a filed one being dropped).

subsequent settlement predictions should be conditioned upon this.

The possibility of NEV suits questions the potential credibility of the plaintiff's trial threat, something which is required to force the defendant to settle. Both Bebchuk (1996) and Nalebuff (1987) suggest ways to make NEV suits credible. Bebchuk does so in a divergent expectations model by recognising that litigation costs tend not to be incurred all at once, and that once they are broken down the plaintiff has smaller 'hurdles' to overcome in order to reach the next stage of the dispute. Alternatively, Nalebuff (1987) amends Bebchuk (1984) by assuming that the uninformed plaintiff's case could be NEV for some types of defendant; this is something the plaintiff can infer from the response to her offer and, as such, may lead her to decide to drop the case. This is of interest for our paper because it takes seriously the possibility of a link between dropping and settlement behaviour, the prospect of the former paradoxically making the latter harder in Nalebuff's model because the plaintiff seeks to avoid 'bad news' by making more aggressive demands.

This brief survey illustrates the variety of models and results on offer when trying to understand the role of case strength in litigation outcomes. In some settings, cases with even chances are most likely to reach trial, in others, more 'extreme' cases (in terms of their strength) make it. Interestingly, few models examine the possibility of cases being dropped once started and the role of case strength on the plaintiff's decision here: high costs, low damages and weak plaintiffs apparently make this most likely, though Nalebuff's model model injects some uncertainty here once the credibility of the litigation threat is taken seriously. Importantly, Nalebuff is rare in treating dropping and settlement as competing outcomes, with the prospect of one affecting that of the other. In turn, this raises an interesting question for models of settlement like Priest and Klein's, where the mechanism that helps the litigation process to select cases for trial is not articulated; presumably, this takes time and could involve the dropping of claims that clearly become weak, as well as the settlement of claims that clearly remain (or become) strong. To the extent that drops and settlements are separate decisions (as modelled within the literature), in Sections 4 and 5 of this paper we investigate their interaction within a competing risks framework.

3 Data

Since 1990, the National Health Service (NHS) in the UK has been decentralised to a significant degree such that individual hospitals have acquired considerable financial autonomy and have adopted commercial accounting practices. Over the same period, moreover, the responsibility for compensating injured patients has shifted first from the individual clinician to the hospital (through a form of enterprise liability), and now to the National Health Service Litigation Authority (NHSLA), a central agency set up to pool litigation risks through what is known as the Clinical Negligence Scheme for Trusts (CNST). The NHSLA has, from April 2002, taken financial responsibility for 100% of all claims against NHS hospitals. However, between 1995 and 2002, under the terms of the CNST, individual hospitals had to retain part of the cost through choosing an 'excess' level, below which they were responsible for the patient's claim. Thus, each hospital during that period chose an excess level under the pooling scheme (the CNST), and this determined the subsequent exposure to liability risk. For many hospitals, therefore, this was a period in which they were financially responsible for a large number of claims under the excess, and hospital claims managers were obliged to liaise with defence and claimant lawyers, as well as claimants themselves, in order to manage the settlement process on behalf of the hospital's board of trustees. In particular, the commercial accounting principles adopted in the 1990's required hospital finance directors to estimate the future liabilities flowing from current claims against the hospital by patients. For this reason, it became important to collect data, both on the financial costs incurred in handling claims, but also on their predicted settlement prospects and expected costs conditional on settlement. To do this accurately, claims managers needed to draw on information supplied by patients (generated during early meetings designed to understand the nature of the patient's claim) and by clinicians or other hospital staff about the alleged negligence. Thus, they were typically in a unique position to make and update estimates of the likelihood that the hospital would lose if the case were to go to trial, and the likely damages payable given the patient's injuries.[12]

Our dataset consists of a full record of all personal injury claims brought against a group of hospitals in one geographic region of the UK, and followed through to their final outcomes—whether they were settled with payment or not. The database was established in the 1980s and was continuously maintained until 2001 and the NHSLA's reforms. Information was collected in relation to the payments made to plaintiffs as well as to defence lawyers at various stages in the settlement process, and the dates of those payments. In addition, as stated above, data were collected in relation to the estimates made at regular intervals after the claim was initiated by the hospitals' claims

[12] In some respects, the claims manager's role was akin to the informal dispute resolution processes described by Farber and White (1991) in their study of US medical malpractice and this may help to explain the similarity with our assessments of case strength identified in the Introduction. Dealing with complaints from an early stage enabled the claims team to diffuse some potential claims while gaining a clear understanding of the basis for others, and their likely strength. Mulcahy *et al.* (1999) cover some of the strategies for handling the transition of complaints into claims in the NHS.

managers about the likelihood that their hospital would be found liable, and, if so, the likely costs and damages. We can have a high degree of confidence in the accuracy of these data because of their contemporaneousness, and their commercial importance in terms of the hospital's risk management, as described above. Data were also recorded on the severity of the patient's injury, and on the estimated source of the patient's funding (state-financed legal aid, trade union, private funds, etc.[13]). Finally, certain procedural events were logged over time by claims managers; in particular, they recorded the dates on which expert reports were received.

As of December 2000 the database contained full information on 3,749 active and closed cases, and full access was obtained for the purposes of the present study.[14] Table 1 shows the descriptive statistics for observed delays and outcomes.

Table 1 here

The mean delay for all claims from incident to initiation was just under three years, with a further delay of around the same duration from the claim's initiation to its closure. For those 1,542 cases which had closed by the end of 2000, approximately 29% (437 claims) had been settled out of court, a further 23 cases (1.5% of the total) went to trial (of which 11 were won by the plaintiff and 12 by the hospital[15]). The remaining 1082 claims (70% of the total) were dropped by agreement or simply abandoned by the patient. For those cases where a payment was made, the average settlement was some

[13]These additional data were collected at our request.

[14]Under the original terms of access, however, we are unable to reveal the identity of the hospitals concerned.

[15]As we saw in Section 2, this approximately equal split in the outcomes of trials in our sample is consistent with the Priest-Klein hypothesis that cases selected for trial are those where the outcome is relatively uncertain, under the assumption of symmetric stakes (see Priest and Klein (1984)).

£17,822, including legal costs. The defence costs incurred were on average £2,209 over all cases, with additional expenses incurred for expert reports and counsel, where needed.

These figures, taken together, illustrate both the 'long tail' nature of litigation against hospitals, and the uncertainty over whether or not a claim will ultimately result in financial liability. They also help to explain the importance of the initial judgements made by claims managers about the likelihood that the patient would win in court, and the subsequent revisions of these that may take place as more evidence becomes available, for example, through expert reports.[16] Table 2 shows the transitions between different estimates made over time by claims managers in respect of liability.

Table 2

In Table 2, the overall effect of these revisions over the course of the litigation process is summarised, showing for each claim the initial and final estimates. Clearly there is a wide range of initial estimates: for some claims the assessment is firmly that the hospital is "not liable", whereas for a similar number of claims it is assessed to be definitely "liable".[17] There are various shades of initial assessments between these extremes. However, as revealed

[16]Sloan and Hoerger (1991) ask a panel of physicians to produce three *ex post* assessments of case strength as the case develops and, as such, they also have updates like ours. Arguably, our data provide more convincing insights because of our larger sample size (Sloan and Hoerger have data from 185 cases), because our updating is not constrained to take place at specific points in the case, and because the contemporaneousness of the assessments means that they are based on all available information—the authors acknowledge the weakness of *ex post* assessments here when discussing the lack of statistical significance in a number of their results.

[17]Sieg (2000) estimates a reasonably tight distribution of liability estimates but both Farber and White (1991) and Sloan and Hoerger (1991) report wide variation. Though not directly comparable, these differences are consistent with Sieg's estimates dating from later in the litigation process than the 'initial' assessments that we, and these other authors, have.

in Table 2, estimates of liability can change over the course of the claim. For example, of the 909 cases that began with an assessment of "unclear" liability, 188 evolved to an estimate that the hospital was definitely "not liable", and 78 to an estimate that the hospital was definitely "liable". Evidently for many claims there is considerable uncertainty initially due to lack of information, and this is gradually resolved over time as more information becomes available about the circumstances of the event and the standard of care adopted by the hospital.

It is important to emphasise the implications of our earlier comments about the ways in which claims managers formulate the estimates in Table 2: these are contemporaneous best guesses derived using information garnered from both sides of the dispute for commercially important reasons. Although the claims managers are working for hospitals it is in their interests to reach accurate assessments of liability and to record them carefully. Table 3 shows that the claims managers' initial assessments of liability are in fact a good predictor of the outcome of settled claims: only 3% of those claims they assess as definitely "not liable" do eventually obtain a paid settlement, whereas almost 88% of those assessed as definitely "liable" obtain a paid settlement. As far as the limited number of trials are concerned, Table 4 shows a similar pattern, with the claims managers' estimates correlated with the actual proportion of cases won at trial (with the overall proportion won being close to 50%).

Table 3 and 4 here

At the same time, claims managers also make an estimate of each claim's value taking into account information provided to them by the patient. These estimates may be revised over time as more information becomes available about the severity of the patient's injury and its prognosis. Table 5 shows that there is a reasonable correlation between the final outcome in terms of damages paid, and the claims managers' assessments of case value, although it does appear that on average, the claims managers are over-cautious and overestimate the likely final settlement.[18]

Table 5 here

In addition, we have data on the timing of the receipt by the hospital of expert reports delivered to them by the patient's representatives. Figure 1 shows the incidence of these reports as litigation time elapses. Most reports were received within 3–18 months of the claim being initiated; the incidence of reports falls subsequently. Typically these are reports by medical experts presenting evidence in relation to standard practice regarding the procedures alleged by the patient to have caused harm, and subsequently made available to both sides. They are a natural source of information for the parties to draw upon in order to update their own assessments of case strength.

Figure 1 here

The data we have described offer a potentially rich insight into the conduct of a large number of medical malpractice claims; not least in terms of the information produced by the litigation process in order to attempt to resolve the case. The remainder of this paper explores the extent to which this information dictates settlement and abandonment probabilities in a predictable way, taking into account measures of severity and the nature of funding used.

[18]See Ayuso and Santolino (2007) for some recent evidence on claims assessment practices in relation to automobile claims. Similar studies have also been undertaken using US insurer data (Gaver and Paterson, 2004; Grace and Leverty, 2007), reporting evidence that insurers have financial incentives to manipulate their overall reserves. For our purposes, we are mainly concerned to show that the relative case value estimates are broadly reflective of relative final payments, and Table 5 appears to support this view.

Other claim level data available to us provide controls for the variations in severity of the patient's injury, and the extent to which the patient's legal costs are borne directly by the patient as opposed to a third party funder. In particular, patients with cases financed by publicly-funded legal aid or trade unions will effectively have their legal costs subsidised.[19] Descriptive statistics for these variables are shown in Table 6.

Table 6 here

4 Estimating methodology

4.1 Cause-specific hazards

In the remainder of this paper we seek to estimate the extent to which the timing of litigation outcomes is determined by the characteristics of the claim as well as the information held by the litigating parties as time elapses. For a given claim, as time elapses and more information about liability and quantum becomes available, this information affects the parties' assessment of their prospects in court and as a consequence the conditional probability of settlement or abandonment for that claim may change over time. Given the nature of our data, we require a regression-based method for estimating the conditional probabilities (hazards) of different types of claim resolution at a given time as a function of relevant covariates. Moreover, as claims may be resolved by one of two competing events—settlement or abandonment[20]—

[19] See Fenn and Rickman (1999).

[20] Note that trial dates are typically set by the courts once the litigation process has begun and, in that sense, their timing is not subject to the litigants' views about the case in any one period—trial is not a competing risk during the litigation process prior to the trial date. For this reason, we treat trials as being censored observations. That is, we assume that they remove the relevant claims from the observation set at the trial

the occurrence of one precludes subsequent observation of the other. Consequently the cause-specific hazards are latent variables. At a given moment in the litigation process, we observe one of three states: either a continuation of the claim, or one of the two outcomes. If we define a latent variable θ_{ijt} as the propensity for claim i to be in state j in Period t, and X_{it} is a vector of covariates (to be defined in Section 5.2), then we can write these propensities as

$$\begin{aligned}
\theta_{i\delta t} &= \beta_\delta X_{it} + \epsilon_{i\delta t} \\
\theta_{i\sigma t} &= \beta_\sigma X_{it} + \epsilon_{i\sigma t} \\
\theta_{i\rho t} &= \beta_\rho X_{it} + \epsilon_{i\rho t}
\end{aligned}$$

where δ, σ and ρ respectively stand for 'abandonment' (the patient unilaterally withdraws her intention to pursue a claim for damages against the hospital), 'settlement' (the patient and hospital agree to end their dispute with an agreed payment of damages) and 'continuation' (the patient decides not to withdraw and the hospital decides not to settle). Over time, as the claim progresses, we could represent the dynamics of the settlement process as being determined by three independent realisations of the random error terms ϵ_{ijt} $(j = \delta, \sigma, \rho)$ from an assumed probability distribution in each discrete time period.[21] In any one discrete period t there are two destination-specific hazards $h_j(t; \beta_j X_{it})$, $j = \delta, \sigma$. Consequently for each claim there are two latent survival times corresponding to each outcome. However, we only observe the minimum of these two—the shortest duration. The other dura-

date—prior to observing whether they would ultimately have settled or dropped had a trial not occurred. As pointed out above, there are only 23 trials in our dataset.

[21]Because of the need to incorporate time varying covariates, we have organised our data into discrete monthly intervals.

tion is latent and unobserved. Note that the assumption of independence of the error terms ϵ_{ijt} ($j = \delta, \sigma, \rho$) is an assumption of *conditional* independence and, as such, is not violated by the latent variables θ_{ijt} being dependent on the same set of covariates. Thus, for example, the propensity to drop a claim and the propensity to settle may both be influenced by the observed measures of case strength (not necessarily in the same direction). The assumption of conditional independence holds that, to the extent these propensities are also influenced by unobservables, the latter act independently across the different means of resolving a claim.[22]

Our data are continuous measurements of litigation durations, which have been rendered discrete by splitting into monthly intervals in order to permit the use of time-varying covariates. In that case, a simplifying assumption might be that the transitions to outcomes take place only at the boundary of the intervals.[23] If so, the assumed functional form for the cause-specific hazards would be the complementary log-log (Prentice and Gloeckler (1978)):

$$h_j(t; \beta_j X_{it}) = 1 - \exp[-\exp(\beta_j X_{it} + \gamma_{jt})] \qquad (2)$$

If we now define λ_{ij} as a cause-specific indicator variable taking the value 1 if claim i is closed with outcome j, and 0 otherwise, it is possible to write the likelihood contribution of an individual claim with an observed duration

[22] While this assumption is conventional in many competing risk contexts (e.g. Narandranathan and Stewart (1993)), it is of course possible that there are unobserved case characteristics which may act dependently across the competing risks. We therefore test for the robustness of our results to the conditional independence assumption using the bivariate probit method suggested by Han and Hausman (1990); the results are reported in Section 5 below.

[23] See Narandranathan and Stewart (1993).

of k periods as follows:

$$L = \left[\frac{h_\sigma(k;\beta_\sigma X_{ik})}{1-h_\sigma(k;\beta_\sigma X_{ik})}\right]^{\lambda_{i\sigma}} S_\sigma(k) \left[\frac{h_\delta(k;\beta_\delta X_{ik})}{1-h_\delta(k;\beta_\delta X_{ik})}\right]^{\lambda_{i\delta}} S_\delta(k) \qquad (3)$$

where $S_j(t)$ is the survivor function for outcome j. The key feature of this expression is that it can be partitioned into the product of two separate cause-specific terms, and therefore the independent competing risks can be estimated separately (using all observations in each case) while still maximising the overall likelihood function.[24] Moreover, the inclusion of monthly dummy variables in the vector of covariates allows a flexible piece-wise linear baseline hazard function to be estimated for each type of outcome. Finally, a parametric specification for the unobserved heterogeneity was suggested by Meyer (1990). In that case the complementary log-log hazard would be specified as:

$$h_j(t;\beta_j X_{it}) = 1 - \exp[-\exp(\beta_j X_{it} + \gamma_{jt} + \ln(u_i)] \qquad (4)$$

where u_i captures the unobserved heterogeneity, and is assumed to have a Gamma distribution with unit mean and variance s^2. The individual effect can then be integrated out by specifying the likelihood function in terms of s^2 as well as the other parameters.[25] We estimate all of the parameters using standard maximum likelihood techniques (Jenkins (1995)). We present estimates for each cause-specific hazard, as well as an overall combined drop

[24]It should be clear from the likelihood function that this approach requires observations on drops to be treated as censored observations in the settlement hazard, and observations on settlements to be treated as censored observations in the drop hazard (see Jenkins (2005); pp 100-101).

[25]Other distributions for the individual effect can be assumed, such as the Normal or the Inverse Gaussian. The estimates we obtained using these alternative distributions were not substantially different. Results are available from the authors.

and settlement hazard to show the conditional likelihood that a claim is resolved during any given period.

5 Results

5.1 Graphical analysis

We begin by graphing the empirical Kaplan-Meier empirical hazards for each of the competing risks of settlement and abandonment, as well as for the overall hazard of the claim being resolved by any means. Figure 2 shows how the outcome-specific hazards change in differing ways as time elapses. The hazard of settlement shows an early rise followed by a period of gradual decline before rising; the drop hazard rises strongly from the early months, becoming relatively much higher than the settlement hazard, only stabilising after several years of litigation time. The overall resolution hazard shows a steady increase throughout.

Figure 2 here

Figures 3 and 4 show the hazards stratified by the claims managers' assessment of case strength. For clarity it was necessary to remove some of the volatility of the hazards from month to month, and this was done in two ways. Figure 3 shows the hazards after smoothing using a weighted kernel function, whereas Figure 4 uses the cumulative hazards as a means of reducing volatility. In both Figures there is a clear difference between the impact of case strength on the hazards of settlement and abandonment. Cases where the claims manager assesses that the hospital is not liable have lower settlement hazards, but higher drop hazards, by comparison with cases where it is assessed as being liable. By contrast, the overall resolution hazard in both

Figures is highest for those cases where claims managers have relatively clear views on liability in either direction (i.e. definitely liable or definitely not liable). The claims with relatively uncertain views over liability are those with the lowest resolution hazards (i.e. they take longest to get resolved). This graphical analysis of the empirical hazards suggests a clear link between the cause-specific hazards and case strength, which we consider in more detail below using regression analysis.

Figures 3 and 4 here

5.2 Regression analysis

Table 7 shows the results from the hazard regressions described above for the settlement, drop and resolution hazards respectively. The discrete time complementary log-log model regression results are reported for each of the competing risks as well as for the overall combined hazard.

The data are organized in person-month form so that time-varying covariates in both models can be included to capture the impact of case value and case strength across cases and over time. One variable (ESTIMATED QUANTUM) measures the recorded assessment of case value for the relevant case in each given time period—that is, in the form of a time-varying covariate. A further set of binary variables (with UNCLEAR liability as the omitted category) measure the claims manager's assessments of the probability that the patient will win in court; again, as explained above, these assessments vary across cases and over time and are included as time-varying covariates.

The arrival of information in the form of expert reports is also captured through time-varying covariates. We are interested in the extent to which this new expert information causes the parties to change their behaviour with respect to dropping or settling the claim. The time-varying covariates used here are therefore presented as a step variable (EXPERT REPORT) taking the value 1 for periods after the receipt of a report, and zero before. Clearly, the arrival of an expert report would be expected to have an effect on the claims manager's assessment of liability, so we also include in each regression a set of interactions between the EXPERT REPORT step variable and the binary variables reflecting the claims managers' assessments.[26]

While we have no direct measure of the per period legal costs faced by the patient, we do have measures of the source of patient finance, which are strong indicators of the extent to which her legal costs are subsidised and these funding types are included in the regressions through dummy variables. Thus, cases with PRIVATE FINANCE are ones funded from the plaintiff's own pocket and, as such, clearly expose her to cost risk.[27] Next, TRADE UNION finance implies the presence of a large funder who is able to co-fund the litigation and, thus, effectively reduce the patient's costs. Finally, a LEGAL AID patient is also shielded from adverse costs in the event of defeat; this time by state-sponsored funding. Our last collection of variables concern the severity of the patient's injuries (measured according to the 'Harvard scales' in HMPS (1990)). We believe that the nature and severity of the patient's injury may have an effect on drop and settlement hazards independently from their effect on damages, if they are associated with uncertainties and delays in assessing the extent of harm.

Table 7 here

For presentational convenience, we have omitted from Table 8 the time

[26]We are grateful to an anonymous referee for this suggestion.

[27]Recall that in England and Wales, the losing litigant typically pays the winner's costs (as well as his/her own). Thus, the prospect of defeat represents a potentially substantial cost risk for the patient.

dummy coefficients which were used to parameterise the flexible piecewise-linear baseline hazards. We have instead used these coefficients to graph the estimated baseline hazards in Figure 5. Comparison with the empirical Kaplan-Meier hazards in Figure 2 is interesting. It is well known that the effect of unobserved heterogeneity potentially biases both the estimated hazard function coefficients and the estimated duration dependence of the baseline hazards (Lancaster, 1990). Consequently, as a comparison of Figure 5 and Figure 2 shows, the reduction in this bias resulting from our use of a Gamma distributed individual effect has resulted in an increase in the rate at which the hazards rise over time, particularly in relation to the drop and resolution hazards. This yields a more plausible 'shape' for the latter hazard, with the conditional likelihood of resolution increasing continually as time elapses.[28]

Figure 5 here

5.2.1 Settlement hazard

The results in relation to the assessment of case strength and case value are statistically significant. Claims where the patient is assessed by the claims manager to have a better chance of prevailing in court (LIABLE or PROBABLY LIABLE) have higher settlement hazards, whereas claims where the hospital is assessed by the claims manager to have a better chance of prevailing in court (NOT LIABLE or PROBABLY NOT LIABLE) have lower settlement hazards. In addition, claims assessed as having a potentially high

[28]While the effect of ignoring unobserved heterogeneity on duration dependence of the baseline hazard can be significant, a recent Monte-Carlo study by Nicoletti and Rondinelli (2010) has shown that its effect on regressor coefficients is negligible, particularly if the baseline hazard is specified flexibly.

settlement value (ESTIMATED QUANTUM) are associated with lower settlement hazards. Each of these are time-varying covariates, and can therefore be interpreted as showing the effect of differences in case strength and value over time as well as over cases. For example, not only do strong cases have higher settlement hazards but, as claims managers change their assessment over time, this is also reflected in the settlement hazard: as cases get stronger, their hazard of settling increases.

As explained above, we also explore the effect of new information on the settlement hazard by considering the arrival of an expert's opinion (EXPERT REPORT). We find strong evidence that the arrival of an expert report is associated with an increase in the settlement hazard in subsequent periods. Note that this result holds after controlling for the direct effect of the claims manager's contemporaneous opinion, so we can infer that it is caused by the parties' reaction to the new information contained in the expert reports. It appears that, as new information arrives from experts, the consequent reduction in uncertainty results in an increased chance of settlement. Thus it seems that, on average, experts bring quicker settlements. The interactions between the arrival of an expert report and the claims managers' assessments of liability are on the whole not significant, although there is an indication that the arrival of a report for a claim where the hospital is believed to be liable makes the claim less likely to settle, whereas the arrival of a report for a claim where the hospital is not believed to be liable makes the claim more likely to settle. That is, although expert reports appear to bring quicker settlements on average, they also make weak cases appear to be stronger and strong cases appear to be weaker. The expert, perhaps predictably, seems to act to correct more extreme views about case strength.

Other variables are also significant. The variables included to capture the effect of plaintiff finance (LEGAL AID, TRADE UNION and PRIVATE) are all positive and significant, with the strongest effects being seen for those cases where the patient funded her own claim. This is consistent with an intuition that higher exposure to costs will induce early settlement.

The coefficients reported in relation to the patient's injury generally support the view expressed above that the greatest delay (lowest settlement hazard) occurs with PERMANENT injuries where the uncertainty surrounding the prognosis is generally higher as a consequence of difficulties in predicting the ultimate impact of the injury on earnings.

5.2.2 Drop hazard

The impact of the claims manager's assessment on the drop hazards reported in Table 8 are statistically highly significant. Claims where the patient is assessed by the claims manager to have a better chance of prevailing in court (LIABLE or PROBABLY LIABLE) have lower drop hazards, whereas claims where the hospital is assessed by the claims manager to have a better chance of prevailing in court (NOT LIABLE or PROBABLY NOT LIABLE) have higher drop hazards. This, of course, is the opposite effect from that found in relation to the settlement hazard which, as we have seen, is higher where the hospital appears to be weak, implying that such cases settle earlier.[29] Claims perceived to be of high value last longer before being dropped, and those funded through LEGAL AID and TRADE UNIONS (and hence with little cost pressure on the patient or her lawyer) tend to have lower drop hazards than those funded privately.

In relation to the effect of EXPERT REPORTS on the drop hazard, the

[29] This finding is also supported by Farber and White (1991) and Sloan and Hoerger (1991).

coefficients are all significantly negative in each of the three drop regressions. On average, patients respond to the news provided by experts by taking a more pessimistic view of their prospects. Again this result holds after controlling for the claims manager's contemporaneous opinion, so we can infer that the change in the drop hazard is caused by the patient's reaction to the new information. The interaction terms between the arrival of an expert report and the liability assessments of claims managers are overall insignificant.

5.2.3 Resolution hazard

The final column of Table 8 shows the regression results for the overall hazard of claim resolution. They show clearly that cases where the hospital is assessed as being NOT LIABLE or LIABLE have the highest cumulative resolution hazards (i.e. they are resolved sooner) while those with more uncertainty in the prognosis (PROBABLY NOT LIABLE, UNCLEAR LIABILITY, PROBABLY LIABLE) are significantly lower. This is consistent with the idea that litigation deals with more 'certain' cases quicker and uses time and the production of information to sort the remaining claims, ultimately leading to drops or settlements in the systematic ways analysed earlier in this section. That is, because cases where the patient is strong are settled early, and cases where the patient is weak are dropped early, it is the cases with unclear or uncertain liability which last longest before being resolved.[30]

[30] Farber and White (1991) present limited evidence in support of this view by looking at the types of cases (and liability assessments) disposed of at summary stages of the dispute. We are able to examine the same issue but from the perspective of the actual durations of the disposition and controlling for changes in assessments of case strength.

5.2.4 Conditional dependence of settlement and drop hazards

As discussed in Section 4, our decomposition of the resolution hazard into separate competing risks of settlement and drop is based on the assumption that these hazards are independent, conditional on the set of covariates used. In order to determine whether our key findings are robust to this assumption being dropped, we follow a suggestion by Han and Hausman (1990) and estimate a bivariate probit regression in which an additional parameter captures the degree of correlation between the competing outcomes, conditional on the covariates (which are parameterized in terms of the bivariate normal distribution). The baseline hazard for each outcome is a flexible piece-wise linear function, estimated as before through the coefficients on monthly dummy variables. For each outcome an error term is defined and the correlation between these is estimated. The results are shown in Table 8.

Table 8 here

The parameter RHO has an estimated value of -0.808, and is highly significant. This suggests that there is indeed a strong degree of conditional dependence between the drop and settlement decisions. There are clearly unobserved characteristics relating to claim or patient which influence drop and settlement decisions in opposite directions. That is, the unobserved heterogeneity terms captured in each of the hazard functions estimated above include factors influencing, for example, perceived case strength, and these may result in a higher probability of settlement but a lower probability of drop. Importantly, however, allowing for this correlation in a bivariate probit regression does not materially affect our findings. Crucially, the relationship between observed case strength and the drop hazard remains negative and highly significant (weak cases are more likely to drop early), whereas the relationship between observed case strength and the settlement hazard remains positive and highly significant (strong cases are more likely to settle early). Thus our core finding (the the most uncertain claims continue for longer because the strong and weak cases are selected out for early settlements or drops respectively) is robust to the presence of conditionally dependent competing risks.

5.3 Discussion

Our results suggest some clear regularities in the ways that assessments of case strength and information appear to affect litigation, and in the ways that the competing risks of dropping and settling interact to identify strong and weak cases over time, leaving the tight ones for trial. It is therefore interesting to consider these findings in the light of some of the models and hypotheses produced by the existing theory, discussed in Section 2. First, in terms of modelling assumptions, the results appear to show that models of litigation explaining *both* drop and settlement behaviour could fruitfully be explored. While some static models of litigation recognise the potential for drops (as we have seen), the requirements imposed by intertemporal rationality have raised problems for incorporating 'drops' in dynamic settings. Since our results also demonstrate the dynamic nature of litigation, with dropping and settling taking place throughout the litigation process, this is a gap for future research.

Continuing the theme of modelling assumptions, we have not sought to comment on whether our data are more consistent with divergent expectations models (where parties assess claim strength on the basis of a common information pool) or asymmetric information (where parties may seek to keep back unfavourable information during settlement negotiations)[31]: we have been interested in whether the process of dropping and settling can help to sort cases on the basis of case strength as assessed by a contemporaneous expert. To the extent that the results are consistent with intertemporal sorting on case strength they indicate that litigation may not resemble a straightforward common values problem as suggested by theory.[32] The role of information that we have observed is analogous to that in the equilibria derived by Cramton (1984) and Admati and Perry (1987). In each of these papers, the offering party designs a sequence of offers to convince the receiver of its beliefs; in particular, stronger offerers use patience to signal their strength. Therefore, these are less likely to settle in any period, in comparison with offerers who have weak beliefs about their positions and who will be more likely to have made an acceptable offer. Although neither of these papers allows for dropping out of the bargaining process, related work by and Cramton and Tracy (1992) considers bargainers (trade unions) with walk-out as well as settlement options. Future research could build on this.

Second, by suggesting that litigation screens strong and weak cases out, leaving intermediate strength ones for trial, our results seem to support the Priest and Klein (1984) selection hypothesis and the version of Friedman and Wittman (2007)'s litigation model that assumes cases are 'high cost'. The results appear less consistent with the one-sided asymmetric information models of Bebchuk (1984) and others, and the prediction for 'low cost' cases in Friedman and Wittman (2007).[33] There are two comments to make here.

[31]Other empirical papers offer a mixed picture here, with Waldfogel (1998) finding support for asymmetric information models and Osborn (1999) favouring divergent expectations. Both studies focus on data from trials.

[32]In such settings, time does not screen case types.

[33]Of course, the definition of a 'high/low' cost case is relative, and although we control for costs in our regressions, this may be within a class of generally high cost cases. Thus,

First, by incorporating both drops and time, we believe our paper provides an answer to the question of *how* litigation might achieve the selection effects that Priest and Klein discuss: it uses and produces information about cases and provides an opportunity for this to be conveyed and tested so that the parties can ultimately achieve a sufficiently common view upon which to act—occasionally, this is one of sufficient uncertainty that trial ensues but more often it is a clear steer towards the payment or abandonment of the claim. Second, this does not mean, in itself, that the asymmetric information models are rejected: the potential addition of 'drops' in such models could be an important development here.[34]

We have suggested that the process of litigation may help produce case outcomes that are consistent with underlying case strength. Although our empirical work is positive rather than normative, it is interesting to consider possible welfare implications from such findings. Friedman and Wickelgren (2010) have recently questioned the welfare effects of settlement in a Bebchuk (1984)-style model of litigation. They note that settlement can reduce welfare because inaccurate litigation outcomes reduce the marginal benefit from engaging in (costly) care and thus dilute the deterrent effects of litigation. This outcome is especially apparent in a one-shot model of litigation where accuracy cannot be improved by learning over time, but a dynamic model of litigation (especially with new information being produced as the case proceeds) would help to mitigate this effect—a point acknowledged by the authors (see p. 153). To the extent that our work demonstrates the ability of litigation to promote 'accuracy' over time, it may also identify some welfare

further work would be interesting to identify objectively high/low cost cases in order to test Friedman and Wittman (2007)'s model.

[34]Though we note that Nalebuff (1987) produces predictions on the effect of case values that are less consistent with our results than Bebchuk (1984)'s. We have seen that Friedman and Wittman (2007) recognise the omission of 'drops' from their model.

gains from this potentially expensive process.

6 Conclusion

This paper has sought to use unique data (drawn from medical malpractice claims in the UK) on the expert assessment of case strength (as it differs across claims and over time) and on reports that can reasonably be expected to have affected the parties' assessment of case strength, in order to estimate the relationship between these assessments and the relative speed with which medical malpractice claims get resolved over time. By estimating this relationship within a competing risk framework we have been able to explore the distinct ways in which the evolution of perceived case strength has affected the decisions to drop and settle respectively. In turn, we have suggested that our results may be interpreted as the litigation process producing information that could help to identify stronger or weaker cases and treat them accordingly. The culmination of this process results in a non-linear relationship between observed case strength and the time to resolution: cases which are resolved soonest are those which are either relatively weak or relatively strong.

Our data demonstrate considerable variation in the initial assessments of case strength and value across claims, as well as the significant evolution of these assessments over the course of litigation. Moreover, these variations across claims and over time are systematically related to the timing of the claim's disposition. We find that lower value claims, and those where the patient's case is assessed to be clearly strong, settle sooner than others; lower value claims, and those where the hospital's case is assessed as clearly strong, are dropped sooner than others. We also find that information produced by expert witnesses during the litigation process has a significant effect on both drop and settlement hazards: it appears that an important function of the litigation process is about producing the information necessary to encourage completion of the case. The information created at the start of litigation and augmented by the arrival of additional expert evidence in relation to the strength of the patient's claim combine to hasten the resolution of strong claims and reduce the likelihood of weak claims receiving a payment. The overall result of these effects is to sort 'certain' cases from 'uncertain' ones and resolve the former sooner, ceteris paribus.

Our results may have broader implications. Delay in medical malpractice litigation can be considerable (and costly), and has attracted attention from those who argue for tort reform in this area. Yet such delay may also play the valuable role of ensuring that the legal system can reach accurate outcomes against a background of considerable uncertainty and our findings are consistent with the tort system providing a means by which well-supported claims are more likely to be compensated. It is therefore important to understand better the underlying processes which determine why some cases are resolved early and others are not. In principle, this is achieved (partly, at least) by the opportunities legal procedure presents for the production and sharing of information and our results complement those of others (Farber and White (1991); Sloan and Hoerger (1991)) in providing some support for the operation of this mechanism.

Clearly, much work remains to be done in order to help us understand litigation and design policies to improve its operation: as well as our empirical results, our paper has pointed to potentially useful areas for theoretical modelling. We feel the current paper's competing risks framework, combined with reliable data on case strength, makes a helpful contribution to the achievement of these goals.

References

Ayuso, M. and Santolino, M. (2007). Predicting automobile claims bodily injury severity with sequential ordered logit models. *Insurance Mathematics and Economics*, **41**(1), 71–83.

Bebchuk, L. A. (1984). Litigation and settlement under imperfect information. *RAND Journal of Economics*, **15**(3), 404–415.

Bebchuk, L. A. (1996). On divisibility and credibility: The effects of the distribution of litigation costs over time on the credibility of threats to sue. Discusion Paper No. 190, 8/96, Harvard Law School, Cambridge MA 0213, US.

Chatterjee, K. and Samuelson, W. (1983). Bargaining under incomplete information. *Operations Research*, **31**(5), 835–851.

Cooter, R. D. and Rubinfeld, D. L. (1994). An economic model of legal discovery. *Journal of Legal Studies*, **23**(1), 435–463.

Danzon, P. M. and Lillard, L. A. (1983). Settlement out of court: The disposition of medical malpractice claims. *Journal of Legal Studies*, **12**(2), 345–377.

Daughety, A. and Reinganum, J. (2008). Settlement. In C. W. Sanchirico, editor, *Procedural Law and Economics*, volume 10 of *Encyclopedia of Law and Economics*. Edward Elgar Publishers, 2nd edition. Forthcoming.

Daughety, A. F. and Reinganum, J. F. (1994). Settlement negotiations with two-sided asymmetric information: Model duality, information distribution, and efficiency. *International Review of Law and Economics*, **14**, 283–298.

Farber, H. and White, M. (1991). Medical malpractice: An empirical examination of the litigation process. *RAND Journal of Economics*, **22**(2), 199–217.

Farber, H. and White, M. (1994). A comparison of formal and informal dispute resolution in medical malpractice. *Journal of Legal Studies*, **23**(2), 777–806.

Farmer, A. and Pecorino, P. (2002). Pretrial bargaining with self-serving bias and asymmetric information. *Journal of Economic Behaviour and Organisation*, **48**, 163–176.

Fenn, P. and Rickman, N. (1999). Delay and settlement in litigation. *Economic Journal*, **109**(457), 476–491.

Fenn, P. and Rickman, N. (2001). Asymmetric information and the settlement of insurance claims: Theory and evidence. *Journal of Risk and Insurance*, **68**(4), 615–630.

Fournier, G. and Zuehlke, T. (1996). The timing of out-of-court settlement. *RAND Journal of Economics*, **27**(2), 310–321.

Fournier, G. M. and Zuehlke, T. W. (1989). Litigation and settlement: An empirical approach. *Review of Economics and Statistics*, **71**(2), 189–195.

Friedman, D. and Wittman, D. (2007). Litigation with symmetric bargaining and two-sided incomplete information. *Journal of Law, Economics, and Organisation*, **23**(1), 98–126.

Friedman, E. and Wickelgren, A. (2010). Chilling, settlement and the accuracy of the legal process. *Journal of Law, Economics, and Organisation*, **26**(1), 144–157.

Gould, J. P. (1973). The economics of legal conflicts. *Journal of Legal Studies*, **2**(2), 279–300.

Han, A. and Hausman, J. A. (1990). Flexible parametric estimation of duration and competing risk models. *Journal of Applied Econometrics*, **5**(1), 1–28.

Helland, E. and Tabarrok, A. (2003). Contingency fees, settlement delay and low-quality litigation: Empirical evidence from two datasets. *Journal of Law, Economics, and Organisation*, **19**(2), 517–542.

HMPS (1990). *Patients, doctors and lawyers. Medical injury, malpractice litigation, and patient compensation in New York: The report of the Harvard Medical Practice Study to the State of New York*. Harvard Medical Practice Study, Harvard University, Cambridge.

Huang, K.-C. (2009). Does discovery promote settlement? An empirical answer. *Journal of Empirical Legal Studies*, **6**(2), 241–278.

Hylton, K. N. and Lin, H. (2009). Trial selection theory and evidence: A review. Working Paper, School of Law, Boston University.

Jenkins, S. P. (1995). Easy estimation methods for discrete-time duration models. *Oxford Bulletin of Economics and Statistics*, **57**(1), 757–782.

Jenkins, S. P. (2005). Survival analysis. `http://www.iser.essex.ac.uk/teaching/degree/stephenj/ec968/pdfs/ec968lnotesv6.pdf`.

Kessler, D. (1996). Institutional couses of delay in the settlement of legal disputes. *Journal of Law, Economics, and Organisation*, **12**(2), 432–460.

Landes, W. M. (1971). An economic analysis of the courts. *Journal of Law and Economics*, **14**(1), 1–2.

Loewenstein, G. and Moore, D. A. (2004). When ignorance is bliss: Information exchange and inefficiency in bargaining. *jls*, **33**(1), 37–58.

Mulcahy, L., Lloyd-Bostock, S., and Rosenthal, M. (1999). *Medical Mishaps: Pieces of the Puzzle*. OUP, Oxford.

Nalebuff, B. (1987). Credible pretrial negotiation. *RAND Journal of Economics*, **18**(2), 198–210.

Narandranathan, W. and Stewart, M. (1993). Modelling the probability of leaving unemployment: Competing risk models with flexible base-line hazards. *Applied Statistics*, **42**, 63–83.

Nicoletti, C. and Rondinelli, C. (2010). The (mis)specification of discrete time duration models with unobserved heterogeneity: A Monte Carlo study. *Journal of Econometrics*, **159**, 1–13.

Osborn, E. (1999). Who should be worried about asymmetric information in litigation? *International Review of Law and Economics*, **19**, 399–409.

Posner, R. S. (1973). An economic approach to legal procedure and judicial administration. *Journal of Legal Studies*, **2**(2), 399–458.

Prentice, R. L. and Gloeckler, L. A. (1978). Regression analysis of grouped survival data with application to breast cancer data. *Biometrics*, **34**(1), 57–67.

Priest, G. L. and Klein, B. (1984). The selection of disputes for litigation. *Journal of Legal Studies*, **13**(1), 1–55.

Reinganum, J. F. and Wilde, L. L. (1986). Settlement, litigation, and the allocation of litigation costs. *RAND Journal of Economics*, **17**(4), 557–566.

Shavell, S. (1982). Suit, settlement, and trial: A theoretical analysis under alternative methods for the allocation of legal costs. *Journal of Legal Studies*, **11**(1), 1–2.

Sieg, H. (2000). Estimating a bargaining model with asymmetric information: Evidence from medical malpractice disputes. *Journal of Political Economy*, **1088**(5), 1006–1021.

Sloan, F. A. and Hoerger, T. J. (1991). Uncertainty, information and resolution of medical malpractice disputes. *Journal of Risk and Uncertainty*, **4**(4), 403–423.

Spier, K. E. (1992). The dynamics of pretrial negotiation. *Review of Economic Studies*, **59**, 93–108.

Spier, K. E. (2007). Litigation. In A. M. Polinsky and S. Shavell, editors, *Handbook of Law and Economics*. Elsevier North Holland.

Studdert, D., Mello, M., Gawande, A., Gandhi, T., Kachalia, A., Yoon, C., Puopolo, A., and Brennan, T. (2006). Claims, errors and compensation patments in medical malpractice litigation. *New England Journal of Medicine*, **354**(19), 2024–2033.

Viscusi, W. K. (1986). The determinants of the disposition of product liability claims and compensation for bodily injury. *Journal of Legal Studies*, **15**(2), 321–346.

Viscusi, W. K. (1988). Product liability litigation with risk aversion. *Journal of Legal Studies*, **17**(1), 101–121.

Waldfogel, J. (1998). Reconciling asymmetric information and divergent expectations theories of litigation. *Journal of Law and Economics*, **41**, 451–476.

Watanabe, Y. (2006). Learning and bargaining in dispute resolution: Theory and evidence from medical maplractice litigation. Mimeo, Northwestern University.

TABLES AND FIGURES

Table 1: Descriptive statistics - outcomes

	Mean	Std. Deviation
Delay from incident to claim (days)	1066	1835
Delay from claim to closure (days)	1004	759
Settled with payment (%)	29.05	44.9
Damages paid to plaintiff inc costs (£)	17822	49530
Defence costs (£)	2209	7957
Payments to expert witnesses (£)	286	1505
Payments to counsel (£)	254	2104

Table 2: Transition matrix – claims

	Final estimate:				
Numbers of claims	Not liable	Probably not liable	Unclear	Probably liable	Liable
Initial estimate:					
Not liable	417	13	15	5	8
Probably not liable	67	212	23	18	25
Unclear	188	158	386	99	78
Probably liable	15	9	15	118	78
Liable	4	2	6	2	194
Total	691	394	445	242	383

Table 3: Settlement outcomes by liability estimates

Initial Liability estimate	Percentage settled with payment %	N
Not liable	3.23	592
Probably not liable	15.77	230
Unclear	20.24	250
Probably liable	55.91	192
Liable	87.58	255
Total	28.77	1519

Table 4: Trial outcomes by liability estimates

Initial Liability estimate	Percentage paid damages %	N
Not liable	25.00	8
Probably not liable	50.00	6
Unclear	62.50	8
Probably liable	100.00	1
Liable	-	0
Total	47.83	23

Table 5: Case value estimates by damages paid

Damages paid (£)	Mean Estimate (£)	N
0-9,999	12,789	279
10,000-19,999	21,869	85
20,000-29,999	33,557	26
30,000-39,999	53,090	16
40,000-49,999	68,976	7
50,000 and over	192,209	23
Total	27,644	436

Table 6: Descriptive statistics – other variables

	Mean
Finance: legal aid	0.4201
Finance: trade union	0.0372
Finance: private	0.0527
Finance: other/not known	0.4900
Injury: temporary minor	0.1533
Injury: temporary major	0.0436
Injury: permanent partial minor	0.1627
Injury: permanent partial major	0.1042
Injury: permanent total	0.0255
Injury: death	0.0534
Injury: other/not known	0.4572

Table 7: Cause-specific complementary log-log hazard regression results (Gamma specification for the unobserved heterogeneity distribution)

	Settle	Drop	Resolve
TIME-VARYING COVARIATES:			
ESTIMATED QUANTUM	-0.0045***	-0.00206***	-0.00260***
	(-3.87)	(-3.40)	(-4.92)
ESTIMATED LIABILITY: NOT LIABLE	-0.804**	1.721***	1.350***
	(-2.74)	(10.09)	(9.81)
ESTIMATED LIABILITY: PROBABLY NOT LIABLE	-0.0779	0.0376	-0.0217
	(-0.31)	(0.27)	(-0.18)
ESTIMATED LIABILITY: PROBABLY LIABLE	0.929***	-1.415***	-0.591***
	(4.53)	(-7.83)	(-4.52)
ESTIMATED LIABILITY: LIABLE	2.574***	-2.453***	0.805***
	(11.74)	(-8.25)	(5.61)
EXPERT REPORT	0.436**	-1.211***	-0.488***
	(3.10)	(-6.83)	(-4.13)
EXPERT REPORT*NOT LIABLE	0.860	0.551	0.132
	(1.46)	(1.48)	(0.48)
EXPERT REPORT*PROBABLY NOT LIABLE	-0.0317	0.541	0.211
	(-0.06)	(1.43)	(0.76)
EXPERT REPORT*PROBABLY LIABLE	-0.674	-0.159	0.328
	(-1.68)	(-0.33)	(1.19)
EXPERT REPORT*LIABLE	-1.016**	0.163	0.434
	(-2.72)	(0.25)	(1.68)
FIXED COVARIATES:			
FINANCE: LEGAL AID	0.437**	0.216	0.302**
	(2.88)	(1.80)	(3.04)
FINANCE: TRADE UNION	1.171***	1.383***	1.194***
	(4.53)	(5.02)	(5.80)
FINANCE: PRIVATE	1.696***	1.444***	1.583***
	(6.98)	(5.85)	(8.11)
INJURY: TEMPORARY MINOR	0.359*	0.0805	0.237
	(2.12)	(0.51)	(1.92)
INJURY: TEMPORARY MAJOR	-0.276	0.128	0.000975
	(-0.90)	(0.48)	(0.00)
INJURY: PERMANENT PARTIAL MINOR	-0.772***	-0.194	-0.349**
	(-3.67)	(-1.25)	(-2.67)
INJURY: PERMANENT PARTIAL MAJOR	-0.836**	-0.0862	-0.340*
	(-3.11)	(-0.43)	(-2.00)
INJURY: PERMANENT TOTAL	-0.0425	-0.500	-0.525
	(-0.07)	(-1.13)	(-1.38)
INJURY: DEATH	-0.0779	-0.0471	-0.0275
	(-0.26)	(-0.20)	(-0.14)
CONSTANT	-4.114***	-1.300***	-1.051**
	(-15.25)	(-3.50)	(-3.00)
GAMMA VARIANCE (s^2)	0.5238***	1.171***	0.8329***
	(3.33)	(5.71)	(5.61)
Obs.	49200	49200	49200
Subjects	1542	1542	1542
Failures	437	1082	1519
Log Likelihood	-2064	-4676	-6361

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

For presentational convenience, we omit from the Table the time dummy coefficients which were used to parameterise the flexible piecewise-linear baseline hazards.

Table 8: Bivariate probit regression results (Han-Hausman dependent competing risks model)

	Settle	Drop
TIME-VARYING COVARIATES:		
ESTIMATED QUANTUM	-0.00137*	-0.000549***
	(-2.39)	(-3.78)
ESTIMATED LIABILITY: NOT LIABLE	-0.321**	0.381***
	(-2.83)	(10.34)
ESTIMATED LIABILITY: PROBABLY NOT LIABLE	-0.0353	0.0115
	(-0.34)	(0.26)
ESTIMATED LIABILITY: PROBABLY LIABLE	0.377***	-0.389***
	(4.64)	(-7.32)
ESTIMATED LIABILITY: LIABLE	0.958***	-0.716***
	(12.67)	(-7.70)
EXPERT REPORT	0.391**	-0.432***
	(2.88)	(-3.39)
EXPERT REPORT*NOT LIABLE	0.266	0.187
	(1.26)	(1.23)
EXPERT REPORT*PROBABLY NOT LIABLE	-0.0272	0.200
	(-0.15)	(1.33)
EXPERT REPORT*PROBABLY LIABLE	-0.216	-0.0370
	(-1.40)	(-0.21)
EXPERT REPORT*LIABLE	-0.318*	0.0834
	(-2.17)	(0.38)
FIXED COVARIATES:		
FINANCE: LEGAL AID	0.124**	0.0555
	(2.59)	(1.81)
FINANCE: TRADE UNION	0.424***	0.359***
	(4.80)	(4.70)
FINANCE: PRIVATE	0.596***	0.337***
	(7.84)	(6.11)
INJURY: TEMPORARY MINOR	0.116*	0.0320
	(2.05)	(0.74)
INJURY: TEMPORARY MAJOR	-0.113	0.0247
	(-1.16)	(0.36)
INJURY: PERMANENT PARTIAL MINOR	-0.266***	-0.0334
	(-4.15)	(-0.83)
INJURY: PERMANENT PARTIAL MAJOR	-0.263***	-0.0416
	(-3.37)	(-0.85)
INJURY: PERMANENT TOTAL	-0.0574	-0.105
	(-0.24)	(-1.10)
INJURY: DEATH	-0.0403	0.0135
	(-0.44)	(0.22)
CONSTANT	-2.338***	-1.800***
	(-27.19)	(-29.77)
RHO	-0.808	
Wald test of rho=0:	295.411	
Pr(rho=0)	0.0000	
Obs.	49200	49200
Subjects	1542	1542
Failures	437	1082
Log Pseudo-likelihood	-6758	
Wald test	1702.66	
Pr(W>Chi²)	0.0000	

t statistics in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Time dummy coefficients omitted for presentational convenience

Figure 1

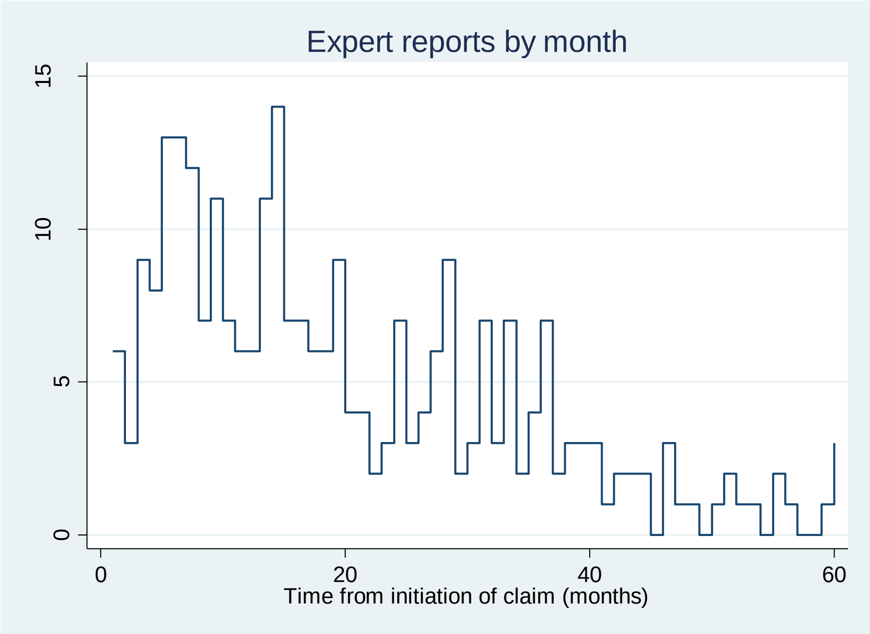

Expert reports by month
15
10
5
0
0
20
40
60
Time from initiation of claim (months)

Figure 2: Kaplan-Meier hazards by type of disposition

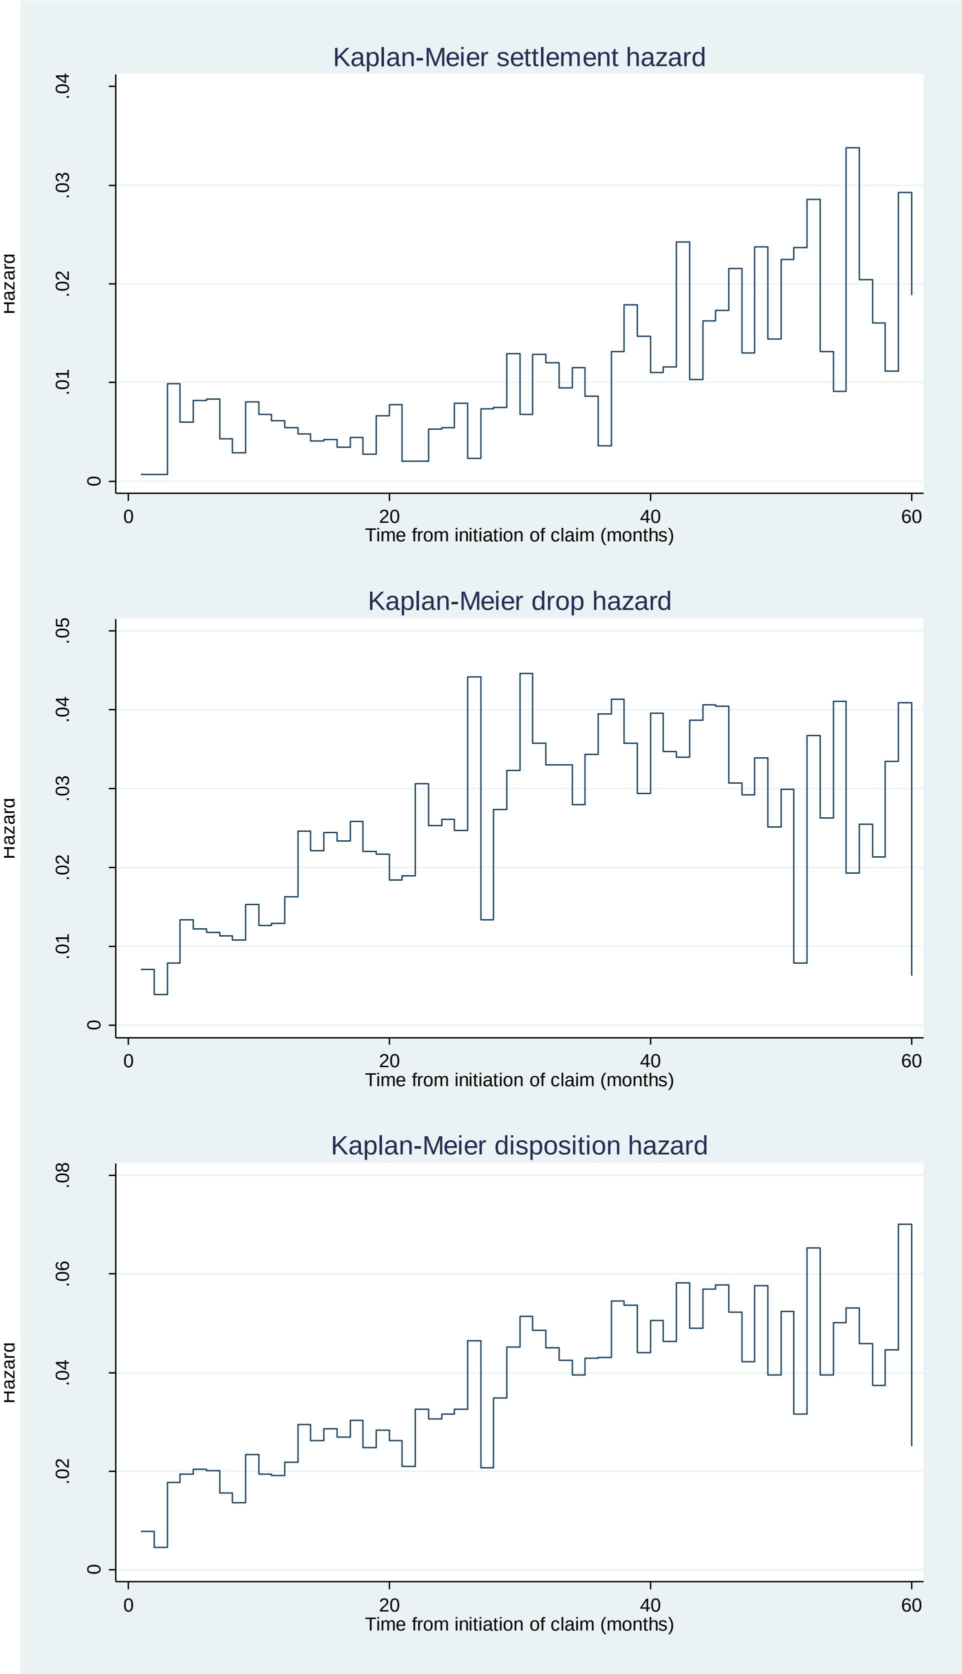

Figure 3: Smoothed hazards by type of disposition and case strength

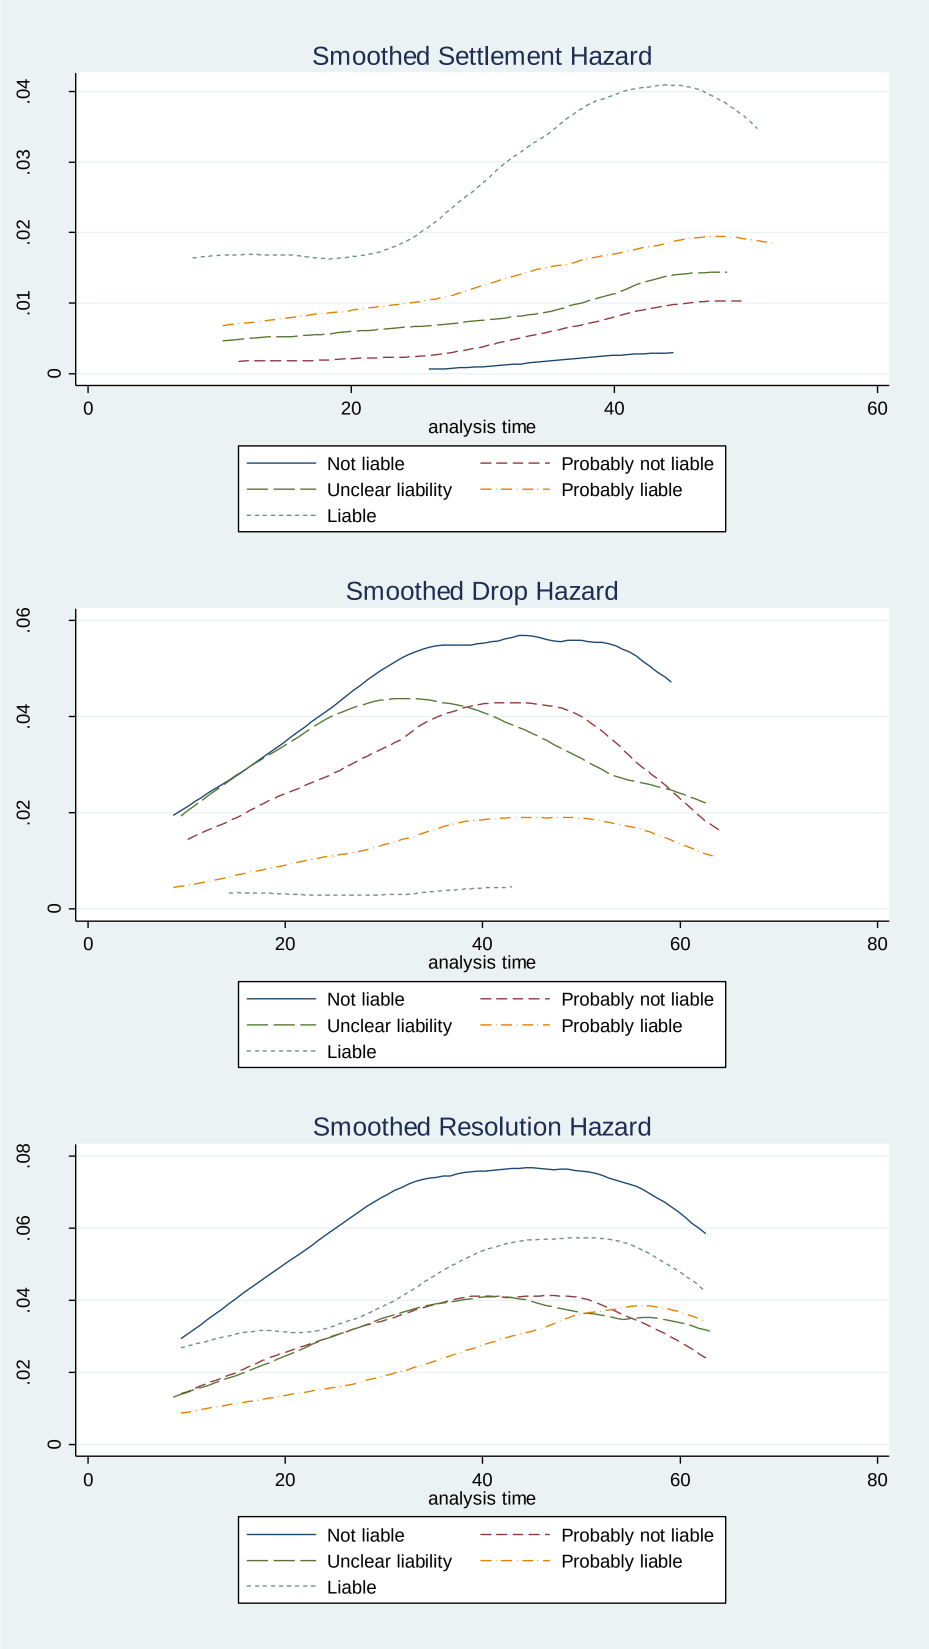

Figure 4: Cumulative hazards by type of disposition and case strength

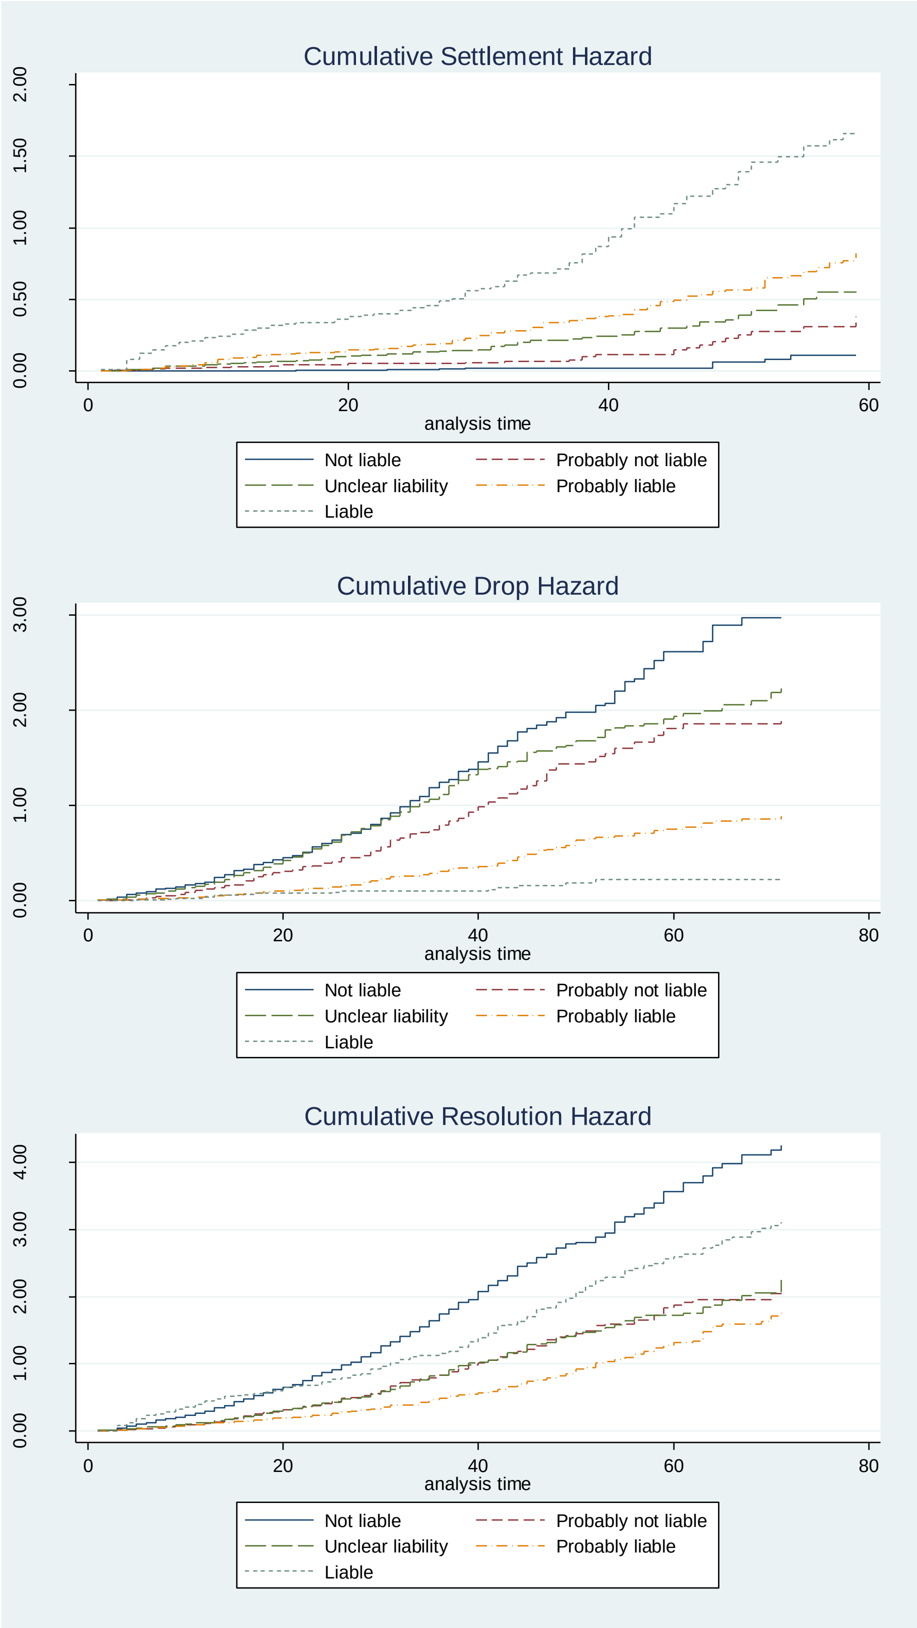

Figure 5: Baseline hazards derived from piecewise linear regressions (with correction for unobserved heterogeneity)