

Reading the Economic Logic of Law from Its Recorded Disputes

The Machloket as Identification

Torun Dewan*

Limor Ressler†

Abstract

A legal rule, even one that states its own rationale, rarely pins down the economic mechanism that generates it. The same rule is usually consistent with more than one account of why it is the rule, so the rule underdetermines the mechanism behind it. Lawyers resolve the ambiguity from outside, asserting which account is correct or treating the variety of rules across systems as harmless. This Article identifies a different resource, one the law supplies from within: its recorded disagreements. In the civil law of the Talmud a rule is often preserved alongside a *machloket*, a dispute among the sages over the rule's reason, together with the subsidiary rulings each disputant accepts. That preserved disagreement, we argue, is an identifying restriction. The rival opinions are rival economic mechanisms; the rule on which they agree is the underdetermined object; and the subsidiary rulings on which they differ are the observations that tell the mechanisms apart. The Article states the method, shows that it is selective rather than universal, and works it through two doctrines, the liability of a renter and the liability for fire, while drawing each beside its common-law cousin. The disagreements a legal tradition records, on this account, are not noise to be smoothed away. They are evidence.

*London School of Economics and Political Science. Draft of June 2026; comments welcome. Talmudic citations follow the standard Vilna foliation and have been checked against Sefaria; translations are our own unless otherwise noted.

†Contributing Lecturer, University of Cambridge, Judge Business School; Member of the New York and Israel Bars. Formerly Barclays Bank plc, Arnold & Porter LLP.

1 Introduction

Every common lawyer knows that an employer answers for the torts of his employee. A delivery driver, in the course of his rounds, runs a red light; the firm pays. The rule is among the most settled in the law of obligations; its reason is among the least settled. The employer may be liable because he controls the employee, whose act is his own at one remove; or because an enterprise that profits from an activity should bear the losses that activity throws off; or because the firm is the party best placed to spread the cost through insurance and price. Each reason yields the rule, and the rule, stated as a result, cannot say which is its own. The disagreement is old enough, and live enough, to carry its own name in the literature: the rival rationales of vicarious liability.¹

One might hope the common law would settle the question by argument. It has not, and the manner of its failure is the instructive part. Where the courts adopt a reason they do not hold the rule fixed and dispute its ground; they change the rule. California, taking the enterprise view, extends liability to harms the employee was in no way engaged to cause; most other jurisdictions, holding to control, deny it.² The rival reasons do not converge on a shared rule and divide over its application. They produce different rules, and different outcomes, in different places. This is the variety that comparative law has long noticed and long discounted, and its most influential statement is Saul Levmore's: where rival rationales generate rival rules, the diversity among legal systems is the harmless residue of behaviourally equivalent law, and a system's selection among reasons does not matter.³ The disagreement tells us the choice of reason is consequential. It does not tell us which reason is right, for it is never staged over a rule the disputants share. But the diversity Levmore discounts is, in the right materials, not harmless

¹Daniel Harris, "The Rival Rationales of Vicarious Liability," 20 *Florida State University Business Review* 49 (2021), distinguishing the agency (control) rationale from the "modern justification" that gathers enterprise liability, loss-spreading, and deterrence.

²Harris, *supra*, at 65: the enterprise rationale "is the law in California, but it is not always followed. Elsewhere, the agency approach is the general rule, subject to occasional exceptions." The divergence is sharpest for an employee's intentional torts; compare *Lisa M. v. Henry Mayo Newhall Memorial Hospital*, 12 Cal. 4th 291 (1995) (no liability under the motive test), with *Ira S. Bushey & Sons v. United States*, 398 F.2d 167 (2d Cir. 1968) (liability on an enterprise-risk rationale). That the agency-based motive test remains dominant across United States jurisdictions, and that it ordinarily bars recovery in workplace sexual-assault cases, is documented in Mark A. Geistfeld, "Reformulating Vicarious Liability in Terms of Basic Tort Doctrine: The Example of Employer Liability for Sexual Assaults in the Workplace," 99 *New York University Law Review* 578 (2024) ("[t]he motive test rules the land").

³Saul Levmore, "Rethinking Comparative Law: Variety and Uniformity in Ancient and Modern Tort Law," 61 *Tulane Law Review* 235 (1986), reading the variety among the tort rules of the ancient codes as economically intelligible and the diversity among them as the harmless residue of behaviourally equivalent rules.

but evidential: where a tradition preserves its disagreements, it records alongside each reason the subsidiary rulings that reason commits to, and those rulings predict different outcomes on specific margins. The choice of reason is then not immaterial but observable.

The Talmud stages it over exactly such a rule. Consider a rule that divides a loss in half. Two neighbours own oxen; one ox gores the other; the law orders the owners to share the loss equally.⁴ Why a half? The rule is consistent with at least two reasons. It might be a penalty, imposed on an owner the law presumes blameless, to make him guard an animal that has shown no vice. Or it might be compensation, owed in principle for a loss one party's property inflicted on another, and softened to a half because the goring could not be foreseen. The two readings agree on the number and disagree on everything behind it: the first treats the baseline as zero and the half as an imposition, the second treats the baseline as full and the half as a mercy. The rule alone cannot tell us which is right. It is, in the language of social sciences, underidentified: the rule is the evidence, and more than one mechanism fits it equally well.

This is the ordinary condition of legal rules. A standard of care, a presumption, a measure of damages, a burden of proof: each is typically the equilibrium of more than one rationale, and the doctrine on its face does not say which of them. The law-and-economics tradition has met this condition in two ways. The first is to supply the reason from outside, on the premise that the law is efficient: one asks what purpose the rule would serve if it were optimal, and reads that purpose back into the rule.⁵ The second we have already met, in the vicarious-liability rule: to treat the variety of rules across legal systems as evidence that the choice among them does not matter, so that where negligence, strict liability, and loss-splitting all produce the same behaviour, a system's selection among them is harmless. Both moves resolve the ambiguity by reasoning about the rule. Neither uses the law's own account of itself.

That a rule does not announce its reason is not a defect of drafting; it is structural. On Hart's account a rule carries an open texture, a settled core and a penumbra the rule itself does not reach, and nothing in the rule supplies the ground on which its extension should be decided.⁶ The opacity that sets our problem is a feature of how rules function, not a flaw to be cured.⁷

⁴Exodus 21:35; Mishnah *Bava Kamma* 1:4. The half-damage of the innocuous ox (*chatzi nezek*) is the locus classicus of a divided liability.

⁵The methodological premise is associated above all with Richard Posner; see Richard A. Posner, *Economic Analysis of Law* (1973 and later editions).

⁶H.L.A. Hart, *The Concept of Law* 124–54 (2d ed. 1994). The settled core of a rule's meaning does not extend to its penumbra, and the rule does not supply the rationale for its own extension.

⁷Lon L. Fuller, *The Morality of Law* (rev. ed. 1969). Fuller's rule-of-law virtues require a rule's

We read the reason from within the law's own record of its disagreement, not from principles the interpreter brings to it; the discriminating material is the tradition's, not the analyst's.

The Talmud offers that account. Its civil law is transmitted not as a code of results but as a record of argument, and where a rule is reported the dispute behind it is reported too.⁸ A rule will be stated, and then two named sages will be set against each other over its reason; and each sage will be shown to accept a set of subsidiary rulings, his answers to smaller related questions, that follow from his reason and not from his rival's. Our claim is that this preserved disagreement does identifying work. The rival opinions are rival economic mechanisms. The rule on which they agree is the object the rule underdetermines. And the subsidiary rulings on which they part are the discriminating observations: the restrictions that, added to the rule, pin down which mechanism generates it. The *machloket* is an identifying restriction.

The Talmud has been read for its economics before, and we should mark how our resource differs. The locus classicus is Aumann and Maschler, who showed that the estate divisions of the Mishnah coincide, case for case, with the nucleolus of an associated coalitional game.⁹ That achievement reads a modern economic object off a *rule*: the division the Mishnah prescribes is the datum, and the work is to find the solution concept it instantiates. Our resource is the opposite face of the law. We read not from what the Talmud decides but from what it disputes; the rule is the agreed object, and the economics is recovered from the disagreement over its reason. The two are complementary. Where Aumann and Maschler find the economics latent in a prescription, we find it in an argument, and the recorded dispute does work that no single rule, however ingeniously decoded, can do.

result to be clear and promulgated, which these results are; they do not require a rule to publish its welfare-economic rationale, so silence about the reason is no defect in the Fullerian sense.

⁸On the dialectical form of the Babylonian Talmud and the preservation of rejected opinions see, classically, the discussion of *eilu va-eilu divrei Elohim chayim*, "these and those are the words of the living God," at *Eruvin* 13b. A defeated opinion is not erased; it is kept, and may be revived. The deeper point, that the Babylonian Talmud is distinctively committed to preserving the *reasoning* behind its rulings, including reasoning it rejects, is David Weiss Halivni's: see David Weiss Halivni, *Midrash, Mishnah, and Gemara: The Jewish Predilection for Justified Law* (1986). That commitment is what makes the subsidiary rulings available to be read; without it the resource this Article exploits would not exist. For the standard academic account of the Jewish legal system, its sources and its modes of reasoning, see Menachem Elon, *Jewish Law: History, Sources, Principles* (Bernard Auerbach & Melvin J. Sykes trans., 1994); and, on the tractates worked here, Jacob Neusner, *Introduction to the Talmud of Babylonia* (1984).

⁹Robert J. Aumann and Michael Maschler, "Game Theoretic Analysis of a Bankruptcy Problem from the Talmud," 36 *Journal of Economic Theory* 195 (1985), reading the estate cases of Mishnah *Ketubot* 93a.

The metaphor is drawn from structural inference, and it is worth making explicit, because it is the organising idea of the Article. A model is underidentified when two configurations of its parameters fit the observed data equally well; one cannot tell them apart from the data in hand. Identification comes from a restriction (an exclusion, an instrument, an additional margin) on which the two configurations imply different observations. Find that margin, observe it, and the ambiguity resolves. A legal rule is the datum that two mechanisms both predict. The subsidiary ruling is the margin on which they predict differently. The Talmud, by preserving both the rule and the subsidiary rulings of each side, supplies the restriction the rule lacks. The contribution of this Article is to name this resource and to show how to use it: a method for reading the economic logic of a legal rule from the disagreements the tradition records about it.

The resource is not the Talmud's alone. Any legal order that keeps its disagreements on the record (the common law in its dissents, European law in the opinions of its Advocates General, a statute in the drafts it discarded) preserves, in principle, the same identifying restrictions. The Talmud is the cleanest specimen, not the only one, and we return to the reach of the method in Part VI. This is, in the end, a paper in legal theory, not in the study of Judaism. The Talmud earns its place because it is the legal corpus that preserves subsidiary rulings alongside the rules they qualify, not because its law is our subject; the method would apply to any tradition that did the same, and Part VI shows what a common-law version looks like.

The method is not universal, and its discipline is the point. Disputes are everywhere in the Talmud; most identify nothing, because they are disputes about facts, about the scope of a category, or about what an earlier authority said. A dispute earns the name *identifying* only when it meets three conditions: the disputants agree on the headline rule; they disagree over which of two economic mechanisms generates it; and each commits to a subsidiary ruling that the two mechanisms predict differently. Stating those conditions, and showing which disputes satisfy them and which do not, is itself part of the contribution. A method that explained every dispute would explain none.

We develop the method through two doctrines, chosen because they exhibit its two characteristic shapes. The first is the liability of a renter for goods in his keeping, where the dispute branches cleanly between two mechanisms and the law must choose one. The second is liability for the spread of fire, where the dispute is not a branch but a nesting: a primary rationale of liability that always applies and a supplementary ground laid over it, and the law, in the end, keeps both. Each doctrine we set beside its common-law analogue: the graded bailment of the common law for the renter, the act-and-omission

distinction and the rule of the dangerous animal for fire. Two further doctrines, the goring ox and the cap on ransom, we treat more briefly, because each is the subject of a companion study; they confirm that the method recurs across the corpus and is not an artefact of the two cases worked here.¹⁰

A word on what is and is not claimed. We do not claim that the sages of the Talmud were economists, or that they reasoned in the terms used here. The method is interpretive: it reads a modern mechanism off an ancient dispute, using the subsidiary rulings the tradition records to discipline the reading. What makes the reading more than projection is that the discriminating observations are the law's, not the analyst's. The sages fixed the subsidiary rulings; the Article only shows what those rulings imply about the reason for the rule. Part II states the method. Parts III and IV work the two doctrines. Part V takes the step we think most novel: where the law later settles a dispute, the settlement itself records which mechanism prevailed, so that a resolved disagreement identifies not only the candidate mechanisms but the one the law came to adopt. Part VI turns to objections and to the method's reach beyond the Talmud.

2 The Method

2.1 Rules underdetermine their reasons

Begin with the condition the method addresses. A legal rule is a result: pay half, pay for a fire's spread, bear the loss for theft. A reason is an account of why that result is the law: to deter, to compensate, to allocate a risk, to spare a cost of proof.¹¹ The relation between the two is many-to-one. A single result is the equilibrium of several reasons, which converge on the result and diverge elsewhere. The half-damage of the goring ox is a penalty or a compensation. The liability for the spread of fire answers to the act of the one who lit it or to the property he failed to guard. The liability of a renter is priced by the benefit he draws from the goods or by the reward he is paid to guard them. In each the result is the same under either reason, and the doctrine, stated as a result, cannot

¹⁰The goring ox is developed in *What the Court Cannot See* (draft, 2026); the ransom cap of *pidyon shvuyim* (*Gittin* 45a) in a companion paper on collective commitment. Both are binary branches of the kind described in Part III.

¹¹The deterrence-and-compensation framing of accident law is Calabresi's; see Guido Calabresi, *The Costs of Accidents* (1970), and Guido Calabresi & A. Douglas Melamed, "Property Rules, Liability Rules, and Inalienability," 85 *Harvard Law Review* 1089 (1972). We borrow the distinction only to name the rival mechanisms a rule may serve; its systematic development across the law of harm we leave aside here.

choose between them.

Confronting this, the analyst ordinarily reasons from the rule to a reason. If the law is taken to be efficient, one asks which reason would make the rule optimal and attributes that reason to the rule.¹² If the law is taken to be various but harmless, one observes that the competing reasons would produce the same conduct and concludes that the choice among them is immaterial. Both are inferences *about* the rule. Neither consults the law's own record of the disagreement, because most legal materials do not contain one: a common-law opinion announces a holding and a single rationale, and the rationales the court rejected survive, if at all, in a dissent. The Talmud is unusual in preserving the disagreement as part of the data.

2.2 The machloket as an identifying restriction

Let a rule R be consistent with two mechanisms, M_1 and M_2 , in the sense that each implies R . On the rule alone the mechanisms are indistinguishable. Suppose, however, that there is some further question A (a subsidiary application, a related case) on which M_1 and M_2 imply different answers. A ruling on A then discriminates: it is consistent with one mechanism and not the other. If the law records a ruling on A , it has, perhaps without intending to, identified the mechanism behind R .

A *machloket* of the right kind records exactly this. Two sages agree on R and disagree on its reason; one reasons by M_1 , the other by M_2 ; and each is shown to rule on the subsidiary question A in the way his reason requires. The dispute is thus a pair of fully specified positions, each pairing a reason with its consequences. To read the mechanism behind R , one need only attend to the subsidiary rulings the sages attach to it. The disagreement is the instrument; the subsidiary rulings are the discriminating observations; the rule is identified.

It helps to borrow the econometrician's distinction between two grades of identification. The rule on its own *set*-identifies the mechanism: it narrows the field to those mechanisms consistent with it, $\{M_1, M_2\}$, but no further. The discriminating subsidiary ruling *point*-identifies: it selects one element of that set. The ruling does this work only if it is fixed independently of R , the way an exclusion restriction must be fixed independently of the equation it identifies; a subsidiary ruling that is merely R restated is collinear with R and adds nothing. The cleanest cases are therefore those in which the discriminating

¹²This is the dominant move; its risk is circularity, since almost any rule can be rationalised as optimal under some specification of costs.

ruling lies *away* from the disputed margin. Where instead it coincides with the very point in dispute, the dispute alone cannot point-identify the mechanism, and the restriction must be supplied by an uncontested feature of the surrounding scheme. We meet both situations below: the renter's discriminator is the disputed margin itself and must be braced by an anchor fixed outside the quarrel, while fire's discriminator is a question apart from the dispute and identifies on its own. So the headline claim is not that the bare disagreement settles everything; it is that the disagreement, read together with the subsidiary rulings the sages commit to and the scheme's uncontested anchors, supplies the restriction the rule lacks.

Three conditions must hold for a dispute to do this work, and we will hold the doctrines below to all three.

First, agreement on the rule. The disputants must share the headline result. If they disagree on the result itself, there is no common object whose reason is in question; there is simply a dispute about what the law is. Identification requires a fixed rule and a contested reason, not a contested rule.

Second, disagreement over the mechanism. The dispute must be over which of two economic mechanisms generates the rule, and not over a point of fact, a scriptural derivation, or the scope of a defined term. Many disputes are of the latter kinds and identify nothing economic.

Third, a discriminating subsidiary ruling. Each side must commit to a ruling on some further question that the two mechanisms decide differently. Without such a ruling the disagreement is bare assertion, two labels for the same result, and carries no information. With it, the mechanism becomes observable.

The conditions are demanding, and most recorded disputes fail the second or the third. That is the method's discipline, not its weakness: it picks out the disputes that carry economic information and sets aside those that do not. What follows shows the conditions met, and shows the two shapes the identification can take.

2.3 The maintained assumption

One assumption underlies the method, and we state it plainly rather than disguise it. We treat each disputant's subsidiary ruling as an independent observation. The sage settles it on its own terms, not by deduction from the headline rule. Were the ruling merely the rule restated, it would carry no identifying force. The two observations would coincide, and the mechanism would stay hidden. We see no way to dispense with this.

Identification must come from somewhere, and here it comes from the dispute.

The assumption is untestable, as the identifying assumption of any such method must be. What earns it is the third condition above: the subsidiary rulings are the sages', not ours. The amoraim fix the subsidiary question apart from the rule. Reading its answer as a distinct datum is faithful to the sugya, not imposed on it. We name the assumption here, once, and rely on it wherever a *machloket* does identifying work below.

2.4 Why the rule must matter

A premise underlies all of this: that the rule does economic work, so that the mechanism behind it is worth recovering. The Coase theorem appears to deny the premise. Where bargaining is costless, the assignment of liability does not affect the efficient outcome; the parties trade their way to it whatever the rule, and the rule's economic logic is idle.¹³ The objection must be met and it is, three times.

First, the settings these doctrines govern are the high-transaction-cost ones, where the benchmark does not hold. A fire crosses to a stranger's field; an ox gores an animal its owner never met; a captor holds a hostage. These are not parties who have bargained, or could; the rule allocates a real loss because no costless trade stands ready to undo it. The one doctrine whose parties do deal with each other, the bailment, is governed by a *default*: the graded tiers apply unless the parties stipulate otherwise, and the law lets them stipulate. But a default is itself an economic choice: it governs whenever the parties are silent, and its design allocates the risk of that silence. The mechanism behind the default is exactly what is in question.¹⁴

Second, the externalities run to parties who are not at the table and cannot be. The ransom paid for one captive falls on the next, who is not yet taken; the activity a fire-setter chooses bears on whoever lies in its path. Coasean bargaining internalises an externality between the two who bargain; it cannot reach a third party, still less a future one. Where the harm lands on the absent, the rule is the only instrument, and its mechanism is the whole of the matter.

Third, and most telling for a Coasean, Coase is not the theorem's prisoner. His lesson, where transaction costs are positive, was that the law should place liability so as to

¹³Ronald H. Coase, "The Problem of Social Cost," 3 *Journal of Law and Economics* 1 (1960). The irrelevance holds in the zero-transaction-cost benchmark; Coase's own subject was the world of positive transaction costs, in which the rule does decide the outcome.

¹⁴On default rules as economic instruments, and on information-forcing or penalty defaults, see Ian Ayres and Robert Gertner, "Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules," 99 *Yale Law Journal* 87 (1989).

economise on them, on the party who can avoid the harm at least cost. That is the very logic these doctrines encode and the *machloket* identifies: the renter bears the loss as the party who controls and benefits, the fire-setter bears the activity he alone chooses. The doctrines are answering the question Coase posed for the world he cared about, and the disagreements we read are disagreements about how to answer it. Far from defeating the method, Coase is its warrant.

3 The Binary Branch: The Liability of a Renter

3.1 The graded bailee and the disputed tier

The law of bailment grades the keeper's liability by the terms on which he holds another's goods.¹⁵ An unpaid keeper (*shomer chinam*), who holds as a favour, answers only for negligence. A borrower (*sho'el*), who has the gratuitous use of the thing, answers for everything, even for an accident he could not prevent. Between them stand the paid keeper (*shomer sachar*) and the renter (*socher*), who answer for negligence and for theft and loss, but not for accident. The traditional rationale is a benefit principle: the more a keeper benefits from his possession of the goods, the more of their risk he bears. The keeper who does a favour bears least; the borrower, who has free use, bears most. The scheme is a single ordered typology of four keepers, and the dispute over the renter is a dispute about the whole of it. Each candidate principle must rank all four, from the unpaid keeper at the floor to the borrower at the ceiling; the two principles agree at three of the four rungs and part only at the renter's. To ask why the renter sits where he does is to ask which principle orders the four.

The renter is the revealing case, because his position is ambiguous on the benefit principle, and the sages divide over it. One opinion classes the renter with the paid keeper: he answers for theft and loss. The other classes him with the unpaid keeper: he answers only for negligence.¹⁶ The result on which they agree is that the renter is a bailee within the graded scheme, liable for negligence and excused for accident. They disagree only on his tier. The dispute satisfies the first condition.

¹⁵The scheme is biblical (Exodus 22:6–14) and is set out in Mishnah *Bava Metzia* 7:8 and Mishnah *Shevuot* 8:1; the Talmudic treatment is at *Bava Metzia* 93a and *Shevuot* 49a.

¹⁶The dispute is reported at *Bava Metzia* 93a; in its plain form Rabbi Meir places the renter with the unpaid keeper and Rabbi Yehuda with the paid keeper. The Gemara also records that Rabba bar Avuh taught the baraita with the attributions reversed, in order to reconcile it with the anonymous Mishnah, which classes the renter with the paid keeper. The reversal scrambles which sage holds which view; it leaves the two positions, and the choice between them, intact.

3.2 Two mechanisms for one scheme

It satisfies the second as well, because the two tiers answer to two different principles of why a keeper's liability is graded at all.

On the first, liability tracks the *benefit a keeper derives from holding the goods*. The renter has the use of the thing; use is a benefit; and the party who draws the benefit, and who controls the goods while he holds them, is the natural bearer of the risk, the least-cost avoider of loss.¹⁷ On this principle the renter, who enjoys the use, belongs with the paid keeper.

On the second, liability tracks not the benefit a keeper takes but the *reward he receives for safekeeping*. The paid keeper answers for theft and loss because he is *paid to guard*; his heightened duty is the price of his fee. The renter, however, does not receive a fee for guarding; he *pays* for the use of the goods. On this principle he is not compensated to keep them safe, and so he belongs with the unpaid keeper, at the floor of the scheme.

These are genuinely different economic accounts of the same graded liability. The first prices risk by benefit drawn and control held; the second prices it by reward paid for an assumed duty of care. They converge on every tier but the renter's, and they diverge there. The dispute over the renter is the point at which the two principles can be told apart.

3.3 The discriminating ruling

The third condition is met by the subsidiary ruling that follows from each tier. To class the renter with the paid keeper is to make him answer for theft and loss, and to limit his exculpatory oath to the case of accident. To class him with the unpaid keeper is to excuse him from theft and loss, and to let his oath extend to them. The liability for theft and loss is thus an observable consequence on which the two mechanisms part: the benefit principle imposes it, the reward principle withholds it. A reader who knew only that the renter answered for theft would know that the law had priced his liability by the benefit he draws, and not by a reward he is not paid. The ruling identifies the mechanism.

A doubt presses here, and the method must meet it rather than slide past it. The renter's liability for theft is the very margin the sages dispute; to read the mechanism off

¹⁷The least-cost-avoider logic is the standard economic account of bailment liability. See Richard A. Posner and Andrew M. Rosenfield, "Impossibility and Related Doctrines in Contract Law: An Economic Analysis," 6 *Journal of Legal Studies* 83 (1977); and, on bailment specifically, Shirley Renner, "Liability in Bailment in Israeli Law: Rethinking Benefit," 31 *Israel Law Review* 587 (1997).

it alone would be to argue in a circle, for the tier and the discriminating ruling are then one and the same thing. The independent observation lies elsewhere in the scheme. The borrower answers even for accident, at the ceiling of the four, and the tradition grounds his liability in the benefit he draws, *kol ha-hana'ah shelo*, all the benefit is his.¹⁸ The benefit principle accounts for this exactly. The reward principle cannot: a borrower is paid nothing to guard, and on a reward-for-safekeeping logic he should sit at the floor with the unpaid keeper, not at the ceiling. The borrower's tier is fixed apart from the renter's dispute, and only the benefit principle survives it. This is the over-identifying restriction the disputed margin cannot supply on its own: the contested tier read against the scheme's uncontested anchor. It is also why the resolution of Part V runs as it does: the law settles the renter on the principle the borrower already attests.

The branch is clean. The renter is classed with the unpaid keeper or with the paid one; the tiers are disjoint and exhaustive; one cannot be at once liable and excused for theft. This is the method's first shape, the *binary branch*: two mutually exclusive mechanisms, a rule each implies, and a subsidiary ruling that selects between them. The goring ox is a branch of the same shape: the half-damage is a deterrent penalty or a softened compensation, and the rule that a fine cannot be imposed on a defendant's own admission discriminates between them.¹⁹ The cap on ransom is another, where the cap deters hostage-taking or corrects a common-pool externality, and the treatment of a captive who redeems himself tells the two apart.

The common law graded bailment in just this way, and divided over just this kind of question. Its older law sorted bailees into three: the bailment for the bailor's sole benefit, for the bailee's sole benefit, and for the mutual benefit of both, requiring slight, great, and ordinary care respectively.²⁰ A renter for hire sits, in that scheme, in the mutual-benefit class, exactly the ambiguous middle the Talmud disputes. The common law settled the middle by a standard of ordinary care; the Talmud preserved the dispute, and in preserving it left us the means to see which principle the grading served. The two settlements are themselves instructive, and they diverge. The Talmud's resolution

¹⁸*Bava Metzia* 95a; the borrower's liability for *ones*, an accident beyond his control, is derived from Exodus 22:13–14 and rationalised as *kol ha-hana'ah shelo*.

¹⁹*Modeh bi-knas patur*: a self-admission exempts from a penalty but not from a compensation. The principle makes the penalty reading and the compensation reading of *chatzi nezek* observationally distinct. See *Bava Kamma* 15a and the companion study cited above.

²⁰The tripartite scheme descends from *Coggs v. Bernard*, 92 Eng. Rep. 107 (K.B. 1703), and was systematised by Joseph Story, *Commentaries on the Law of Bailments* (1832). Modern courts have largely collapsed it into a unitary reasonable-care standard, but the benefit-graded structure is the same one the Talmud states.

chose within the grading, affirming the benefit principle; the common law's collapse into a single standard of reasonable care abandoned the grading altogether, treating the choice of principle as one that need not be made, the harmless-variety verdict Levmore describes. Where the one tradition decided the reason, the other decided that the reason did not matter. Both decisions are data, and Part V reads them as such.

4 The Nested Augmentation: Liability for Fire

4.1 His arrows or his property

Liability for fire takes a different shape, and shows the method's second form. One who lights a fire that spreads and damages is liable; fire is among the established heads of damage.²¹ Why he is liable is disputed, and the dispute is reported at *Bava Kamma* 22a. Rabbi Yochanan holds that his fire is liable *because of his arrows* (*mishum chitzav*): the fire is an extension of his own force, as if he had shot the harm, and he is its direct cause. Reish Lakish holds that his fire is liable *because of his property* (*mishum mamono*): the fire is a possession of his, like his ox or his pit, and he answers for the damage it does as he answers for any unguarded thing he owns.²² The two agree that the setter is liable. They disagree on the mechanism of his liability. The first two conditions are met, and met cleanly: the act-and-cause account and the property-and-custody account are as sharp a contrast as the law of damage affords.

4.2 Activity and care

The contrast is, in economic terms, the contrast between *activity* and *care*, and fire is the doctrine in which both margins bind at once.²³ Lighting a fire is a deliberate act, and the choice whether and how large a fire to light is an activity-level decision. Containing it is a matter of care. The two mechanisms price the two margins. A boundary should be

²¹Exodus 22:5; *esh* is one of the four *avot nezikin*, the primary categories of damage, in Mishnah *Bava Kamma* 1:1. On the Mishpatim of Exodus 21–22 read with the tools of comparative legal history, see Bernard S. Jackson, *Wisdom-Laws: A Study of the Mishpatim of Exodus 21:1–22:16* (2006), and Bernard S. Jackson, *Essays in Jewish and Comparative Legal History* (1975).

²²*Bava Kamma* 22a. Rabbi Yochanan: the analogy to property fails because property is tangible and fire is not. Reish Lakish: the analogy to an arrow fails because fire does not travel by the actor's own force but is carried by the wind.

²³The distinction is Steven Shavell's; see Steven Shavell, "Strict Liability versus Negligence," 9 *Journal of Legal Studies* 1 (1980), and, for the synthesis, Steven Shavell, *Foundations of Economic Analysis of Law* (2004). A care-only standard induces efficient precaution but leaves the actor's activity level unpriced, and so excessive; full liability for the harm an activity causes prices the activity itself.

marked, as it was for the renter. The labels *chitzav* and *mamono* are the sages'; reading them as the activity margin and the care margin is ours. What keeps the reading from projection is that the ruling that tests it, the payment for personal injury, was fixed by the sages and not by us.

The property mechanism prices care. To answer for one's fire as for an unguarded possession is to bear the harm one failed to contain; it induces the keeper to fence and to watch. But, like a standard of care, it leaves the activity unpriced: the setter who has taken proper precautions is not charged for the inherent hazard of having lit a fire at all, and so lights too freely. The arrow mechanism prices the activity. To treat the setter as the direct cause is to charge him with all the harm his fire reaches, precaution or no; it prices the act of lighting, and so deters the fire that should not have been lit.

This is why the Talmud, having posed the two mechanisms as rivals, does not in the end choose between them. Pressed by a subsidiary difficulty, it concludes that the sage who holds the arrow mechanism holds the property mechanism as well.²⁴ The structure that results is not a branch but a nesting: a primary rationale of custodial liability, and over it a supplementary ground of liability for the act. Fire requires both because a fire is both a deliberate act, whose activity level must be priced, and a containable hazard, whose care level must be priced. To get both margins right the law keeps both grounds.

This also explains why fire, alone among the heads of damage, carries the arrow ground at all. The goring ox is pure custody: the animal acts, and the owner is liable only for having failed to guard it. There is no human act to price, and so no arrow ground, only the property ground of the ox. Fire is the case in which the human himself does the dangerous act, and so the case in which the activity margin, and the act-based ground that prices it, come into play. Fire is the act-based sibling of the goring ox.

4.3 The discriminating ruling, and two limits

What tells the two grounds apart is the treatment of personal injury. If the setter is liable as a direct actor, his liability reaches all the foreseeable harm his fire does, including injury to the person, for which the law of bodily harm awards its several heads of damages. If he is liable only as the keeper of a dangerous possession, his liability is scoped to property, and does not reach the heads of bodily damages. The arrow mechanism pays them;

²⁴*Bava Kamma* 23a: *de-it leih mishum chitzav, it leih nami mishum mamono*, one who holds the arrow reason holds the property reason too. The conclusion is forced by the case of concealed objects, discussed below.

the property mechanism does not.²⁵ The personal-injury ruling is the discriminating observation, and it identifies which ground is in play. A second case is sometimes read the same way, though less cleanly. When a dog carries a burning coal to a haystack, all agree the owner pays in full for the coal's resting place; the liability for the *spread* beyond it is treated variously in the sources, and on one reading it tracks the arrow ground, the setter answering for his own scattered force.²⁶

The shape here differs from the renter's. The renter's mechanisms were exclusive, and the ruling selected one. Fire's mechanisms are nested, a primary rationale and a supplementary ground, and the ruling identifies whether the added ground is present, not which of two regimes governs. This is the method's second form, the *nested augmentation*. It is a looser identification than the branch, and we do not pretend otherwise: the cleanest cases are the branches, and fire earns its place by showing that the method extends to a structure the branches do not capture.

Yet on a second axis fire is the cleaner of the two. Its discriminating ruling, the payment for personal injury, is a question apart from the dispute over why the setter is liable at all, so the observation is genuinely independent of the disagreement it resolves. The renter's discriminating margin, by contrast, is the disputed tier itself, and its identification had to be braced by an observation drawn from elsewhere in the scheme. Structure and independence pull in opposite directions: the renter is the cleaner branch, fire the cleaner instrument.

Two further rulings confirm the economic reading, and both are the ordinary economics of causation. Liability for fire stops at the harm the setter could foresee: damage carried by an extraordinary wind, a *ruach she-eina metzuva* or wind not to be expected, he does not pay (*Bava Kamma* 60a). The freak wind is the boundary of proximate cause, and charging the setter for harm he could not anticipate would over-deter the activity without improving his care. This exemption is settled, and the authorities agree on it: an unusual wind is a species of *oness*, force majeure, and the setter goes free whether his liability is read through the arrow ground or the property one. What the rishonim dispute is its rationale and not its result, the proximate-cause limit on the setter's own force set against ordinary *oness* on a possession; the exemption holds under either reading. It

²⁵*Bava Kamma* 23a yields this as a practical difference: whether the setter pays the additional heads of indemnity when his fire injures a person. (The exemption of concealed objects, *tamun*, is the other.) It is the discriminating observation on which the two mechanisms part.

²⁶*Bava Kamma* 21b–23a. The case is entangled with the separate law of *tzeroros* (damage by indirect force, which itself carries half-damages), so we do not lean on it; the personal-injury ruling above is the clean discriminator, and we note the coal case only as suggestive.

therefore does not tell the two grounds apart, as the payment for personal injury did, but confirms a limit on causation common to both. And concealed objects burned in the field he does not pay for; the law exempts the hidden.²⁷ The exemption places the risk of hidden value on the party who hid it, the one who alone knew it was there and could have placed it elsewhere; it is a limit on liability for a harm whose magnitude the injurer could not foresee.

4.4 The common-law cousins

The common law reached for the same distinctions. Its law of animals graded the keeper's liability by what he knew: an owner was not liable for his beast's first foray but became liable once it had shown its vice and he had reason to know it.²⁸ Its older law of trespass distinguished harm done by a man's direct force from harm done by a thing in his keeping, the very line between the arrow and the property; and Roman law let the owner of an offending animal or slave surrender it in lieu of damages, a cap on liability at the value of the offending thing that is the civilian cousin of the ox's own *mi-gufo*.²⁹ The same split survives in the modern law as the line between misfeasance and nonfeasance, between authoring a harm and failing to prevent one. The *chitzav* reason is misfeasance: the setter's own force does the damage. The *mamono* reason is the narrower duty the common law lays on one who controls a source of danger, to contain it, the duty that grounds liability for an omission in the exceptional case where control creates it.³⁰ Closest of all is the common law's own rule of strict liability for the escape of a dangerous thing. *Rylands v. Fletcher* holds the occupier who brings onto his land something likely to do mischief if it escapes liable for the harm its escape does, and the fire dispute is, in effect, the *Rylands* question asked one level down: does the liability run to the act of kindling or to the ownership of the dangerous condition?³¹ And the modern economics of fire and

²⁷The exemption of *tamun*, concealed property, is the difficulty that forces the nesting at *Bava Kamma* 23a.

²⁸The scienter or "one-bite" rule; the owner answers once he knew or had reason to know of the animal's dangerous propensity. The structure is the goring ox's *tam* and *mu'ad*, and is treated in the companion study.

²⁹On noxal surrender (*noxae deditio*) and its kinship with the body-of-the-ox cap, see Levmore, *supra*. The act-and-thing distinction is the deep structure of the forms of trespass and case.

³⁰The leading instance is *Home Office v. Dorset Yacht Co. Ltd.* [1970] AC 1004, holding the keeper of dangerous persons liable for failing to restrain them; liability for nonfeasance arises from control over the source of the harm, as the *mamono* reason makes the fire's owner answer for the possession he failed to guard.

³¹*Rylands v. Fletcher* (1868) LR 3 HL 330, aff'g *Fletcher v. Rylands* (1866) LR 1 Ex. 265. Blackburn J.'s formulation, liability for "anything likely to do mischief if it escapes," states the property ground;

of proximate cause turns on exactly the activity-and-care and the foreseeability margins the sages divided over. The point is not that the systems borrowed from one another. It is that the same economic problems generated the same distinctions, and that the Talmud, by disputing them aloud, left the clearest record of what the distinctions are for.

5 What the Resolution Reveals

The method so far reads the mechanism off a live dispute. But many of these disputes were, in time, resolved; the law settled on one tier, one reading, one rule. Far from defeating the method, the resolution extends it. Where the law settles a *machloket*, the settlement records which mechanism prevailed.

The renter's case was settled. The law follows the opinion that classes him with the paid keeper: he answers for theft and loss.³² In the terms of Part III, the benefit principle won: the law priced the renter's liability by the use he draws, not by a reward he is not paid. Fire's case was settled the other way: not by a choice between the grounds but by keeping both, the conclusion that the holder of the arrow reason holds the property reason too. There the law judged that both margins, activity and care, had to be priced, and built a liability that priced each.

So the identification has a sequel. The dispute identifies the candidate mechanisms; the later ruling identifies the one the law adopted. A tradition that argues its rules in the open, and then settles them, leaves a double record: of the mechanisms a rule might serve, and of the mechanism it came to serve. The first is read from the disagreement, the second from its resolution.

the act of bringing the thing onto the land, and accumulating it there, is the activity the arrow ground would price.

³²The ruling of the anonymous Mishnah, codified at *Shulchan Arukh*, *Choshen Mishpat* 303 (the paid keeper, and with him the renter, answers for theft and loss and is excused for accident) within the law of bailees that opens at 291, the renter's hire being treated at 307; and by Maimonides at *Mishneh Torah*, *Hilkhoh Sekhirut* 1:1–3, which opens by naming the four keepers and at once governs the renter by the paid keeper's rule (the borrower and the unpaid keeper are treated in the companion *Hilkhoh She'eilah u-Fikkadon*). That the codifiers reach the result by the benefit principle, and not by the reward principle, is the settlement Part V reads.

6 Objections, and the Reach of the Method

6.1 Objections

That disputes are everywhere, and the method cherry-picks. The Talmud disputes almost everything, and if any dispute could be read as identifying a mechanism the method would prove anything. The answer is the three-condition test of Part II, which most disputes fail. A dispute about a fact, about the scope of a category, or about what an earlier sage held, identifies no mechanism; only a dispute over the reason for an agreed rule, carrying a subsidiary ruling on which two mechanisms differ, does. The taxonomy of Parts III and IV (branch and augmentation) is offered as evidence that the disputes which pass the test fall into a structure, rather than being assembled one by one.

It helps to see the test turn candidates away. *Bava Kamma* offers three that tempt and fail. The pit is liable for an animal that falls into it but not for the vessels; one is tempted to read a mechanism into the exemption, a risk of hidden value placed on the party who left it. But the exemption is drawn from the verse (an ox and not a man, a donkey and not vessels; *Bava Kamma* 28b), and what the sages dispute is the derivation, not an economic reason. It fails the second condition. Tooth and foot do no actionable damage in the public domain (Mishnah *Bava Kamma* 2:2, 19b), which looks like an assumption of risk by one who leaves his goods in the road; but the exemption is read from the verse, “and it shall consume in another’s field” (Exodus 22:4), and the disagreement is over the reading. It fails the second condition as well. The colliding carriers (the man with a beam and the man with a barrel, the runner and the walker; *Bava Kamma* 31a–32a) divide a loss between two who each had leave to be there, and the cases look like a comparative-negligence rule awaiting identification. But the sages decide them by distinguishing the facts, who held the right of way and who moved abnormally, not by disputing the reason for a shared rule; there is no agreed rule whose reason is contested, only a contest over the facts, and the case fails at the first condition, before the third is ever reached. Each of the three carries an economic flavour and none carries an identifying dispute. The test that admits the renter and the fire turns these away, which is what keeps the method from proving whatever it likes.

The three that tempt and fail do so at the first or second condition; the third can bind on its own. A dispute may agree on the rule and genuinely divide over its reason, and still identify nothing, because the rival reasons command the same rulings everywhere one can look. The Talmud has its own name for the test. *Mai beineihu* (what is between them?)

is the demand that a dispute over reasons produce a case it decides differently; where the question returns nothing, the disagreement is real but silent, and the method must pass it by. The third condition is that demand, made a requirement of identification.

That the reading is the analyst's, not the sages'. We have conceded this and restate it as the method's boundary. The mechanism is a modern reconstruction; the sages did not speak of activity levels or least-cost avoiders. What disciplines the reconstruction is that the discriminating observations are not invented for the occasion. The sages fixed the subsidiary rulings (the liability for theft, the payment for personal injury), and those rulings are what the mechanisms must match. The method reads a reason off the law's own consequences, not off the analyst's sense of what would be efficient.

That the dispute is the redactor's construction, not the disputants' record. The most searching objection comes not from economics but from the academic study of the Talmud. The anonymous voice that frames a sugya, the *stam*, is no neutral reporter: it supplies reasons the named sages never stated, sharpens their positions into a cleaner opposition than they held, and arranges the discriminating case, the *mai beineihu*, to make the disagreement legible. The source-critical tradition has made this its central finding.³³ Our own doctrines wear the marks of it. The attributions in the renter dispute are reversed, on the Gemara's own account, to reconcile a *baraita* with the anonymous Mishnah; and the conclusion that one who holds the arrow reason holds the property reason too is forced by the case of concealed objects, not asserted by any named *amora*. If the discriminating ruling is the redactor's device, engineered to render a dispute sharp, then it is not the independent observation the method requires, and the identification is rigged by the hand that records it.

We concede the activity of the *stam* and deny the conclusion. The method does not need the discriminating ruling to be the historical utterance of a named sage; it needs the ruling to be fixed independently of the rule whose reason is in question. A redactor who must construct the discriminating case constructs it because the rule alone does not decide it; the construction is evidence that the question is open, not that the answer is forced. And whoever fixes the subsidiary ruling, the law then carries it downstream into its codes: our reading is of the rationale to which the law as a system is committed, not

³³The foundational statements are David Weiss Halivni's account of the anonymous, reason-supplying *stam* (see *The Formation of the Babylonian Talmud*, trans. and ed. Jeffrey L. Rubenstein (Oxford University Press, 2013)) and Shamma Friedman's stratigraphic method separating the redactional layer from its *amoraic* sources (the methodological introduction to his study of the tenth chapter of *Yevamot*, 1977). For an English synthesis see Jeffrey L. Rubenstein, *The Culture of the Babylonian Talmud* (Johns Hopkins University Press, 2003).

of the private intention of a sage, and the *stam* is itself a jurist whose *mai beineihu* is a claim about what the rival reasons entail. The safeguard against a ruling built to fit is the one the renter already showed. Where a discriminating ruling is suspect, we ask whether an uncontested feature of the scheme, settled outside the dispute, corroborates it; the borrower's liability, fixed apart from the renter's quarrel, is what confirms the benefit principle there. An engineered *mai beineihu* cannot recruit the anchors the redactor did not touch. Where no such corroboration exists, and the discriminating ruling stands only as the *stam*'s own, the identification is weaker, and we mark it as such rather than claim it.

That a resolved dispute is a counterfactual. Where the law has chosen, the losing mechanism is, in a sense, hypothetical. But Part V turns this into information: the choice is itself a datum, and the losing mechanism is the foil against which the chosen one is defined. A rule whose rival was rejected is better understood, not worse, for our knowing what it was rejected in favour of.

6.2 Beyond the Talmud

The resource the method exploits is not peculiar to the Talmud in principle, though in practice the Talmud is very nearly the only place one finds it intact. Any tradition that argued its rules in the open, and kept the argument, could in principle preserve the same identifying restrictions. The common law is the natural place to look, and the example with which we began shows what one finds there instead. The rule of vicarious liability is settled; its reason is disputed; but the disputants do not share the rule. Where a court adopts the enterprise rationale it widens the rule, and where it holds to control it does not, so the reasons are paired not with subsidiary rulings under one rule but with different rules in different jurisdictions. That is variety, not identification: it tells us the choice of reason has consequences, not which reason a shared rule encodes. And where the rival rationales are set out in their fullest form, it is usually a treatise or a Restatement that sets them out, reconstructing the reasons after the fact; the discriminating cases are then the commentator's selection, not a commitment the deciding court recorded as its own.

Yet the common law can, on occasion, take the identifying shape, and saying when it does sharpens the method rather than confining it to the Talmud. The structure appears when a circuit split divides not over a rule but over its reason. Suppose the federal courts of appeals agree on a doctrinal rule, divide over the rationale that supports it, and each circuit, holding to its rationale, commits to a subsidiary holding on a related

question that the rival rationale would decide the other way. That is a *machloket* in the common-law key: an agreed rule, two reasons, and discriminating subsidiary rulings each side has placed on the record as its own. The circuits that take one reason rule one way on the satellite question; those that take the other rule the other way; and the split between them identifies, exactly as a *machloket* does, which mechanism each circuit's rule encodes. The difference from the vicarious-liability example is the difference that matters: there the courts changed the rule along with the reason, while here they hold the rule fixed and let the subsidiary rulings carry the disagreement. Such splits are not the common run, for an appellate court more often states one rationale and leaves the rest unspoken; but where they occur they are identifying disagreements of just the kind this Article reads, and they show the resource to be structural, not the property of one tradition.

One example comes about as close as the common law seems to come. All federal courts accept the inevitable-discovery exception to the exclusionary rule: evidence obtained unlawfully is admitted if it would inevitably have been found by lawful means.³⁴ The rule is shared; its reason is not. One line of courts grounds the exception in deterrence alone: suppression is justified only so far as it deters unlawful searches, and where the evidence would have surfaced anyway there is nothing to deter, so it comes in. Another line adds a concern for judicial integrity, that a court must not reward officers who take an unlawful short cut and reconstruct a lawful path after the fact. The two reasons agree on the rule and part on a subsidiary question: may the prosecution invoke the exception when the lawful route had not yet been begun at the moment of the unlawful search? On the deterrence reason the answer is yes, for the hypothetical lawful discovery suffices. On the integrity reason the answer is no, and these courts require that the lawful means were already being *actively pursued* when the violation occurred.³⁵ The active-pursuit holding is the discriminating subsidiary ruling: it follows from the integrity reason and not from the deterrence reason, and a reader who saw only that a circuit required active pursuit would know which reason that circuit's rule encodes.

The fit is not perfect, and the imperfection is instructive. The choice of reason here is partly a choice about the rule's own scope, so the circuits that divide over the reason can be described, under a little pressure, as holding slightly different rules. This is the

³⁴*Nix v. Williams*, 467 U.S. 431 (1984).

³⁵The active-pursuit requirement is stated in *United States v. Satterfield*, 743 F.2d 827 (11th Cir. 1984), and applied in *United States v. Cabassa*, 62 F.3d 470 (2d Cir. 1995), and *United States v. Souza*, 223 F.3d 1197 (10th Cir. 2000); the deterrence-only reading admits the evidence whether or not the lawful investigation was already under way.

common law's standing difficulty, the one the bailment already showed: rule and reason are not cleanly separable when the court that adopts a reason may restate the rule in the same breath. A perfectly clean common-law *machloket*, agreed rule and contested reason and a discriminating subsidiary ruling all held apart, is correspondingly rare. That rarity is not an embarrassment to the method but a confirmation of its point: the Talmud is valuable precisely because it holds the three apart as a matter of form, on the page, where the common law lets them run together.

What the Talmud supplies, and the common law for the most part does not, is the whole structure on a single page: the agreed rule, the rival reasons, and each sage's own subsidiary rulings, fixed by the disputants and not by a later hand. This is why the method finds its cleanest specimen here. The claim is not that only the Talmud could yield identifying disagreements, but that the Talmud is where they have been kept intact, in a civil law transmitted as argument, with the rejected reasons and their consequences left on the record. Wherever else a tradition argues with itself in the open and preserves both the reasons and the rulings each commits to, the same method will apply. The disagreements a legal tradition keeps are an evidentiary resource, and this Article has tried to show how to spend it.

7 Conclusion

A legal rule, even one that states its own rationale, rarely pins down the economic mechanism that generates it. The same rule serves more than one economic purpose, and the rule alone cannot say which. The Talmud, alone among the great legal corpora in the fullness with which it does so, preserves its disagreements over the reasons for its rules, together with the subsidiary rulings each disagreement entails. Those preserved disagreements are identifying restrictions. The rival opinions are rival mechanisms; the agreed rule is the object they underdetermine; the subsidiary rulings are the observations that tell them apart. The renter's liability is priced by benefit or by reward, and the liability for theft tells us which. The liability for fire prices the act or the custody or both, and the payment for personal injury tells us which grounds are at work. Where the law later settles such a dispute, the settlement records the mechanism that won. Read this way, the Talmud's arguments are not obstacles to stating its law. They are the evidence from which its law's economic logic can be recovered.

A The Identification, Stated Formally

For the reader who wants the metaphor made precise. Let a rule R be the observable, and let M_1 and M_2 be two mechanisms, each a mapping from primitives to legal consequences, with $M_1 \Rightarrow R$ and $M_2 \Rightarrow R$: on the rule alone the mechanisms are observationally equivalent, and R does not identify the mechanism. Let A be a further question on which the mechanisms diverge, $M_1 \Rightarrow a_1$ and $M_2 \Rightarrow a_2$ with $a_1 \neq a_2$. A ruling $A = a_i$ is then a discriminating observation: the pair $(R, A = a_1)$ is consistent with M_1 and not M_2 , and conversely. In the language of partial identification, R alone *set*-identifies the mechanism, with identified set $\{M_1, M_2\}$, while the ruling $A = a_i$ *point*-identifies it. The identification maintains that the ruling on A is observed in its own right, fixed independently of R ; this is the exclusion restriction that licenses the second step. A ruling deducible from R is collinear with it: it leaves the set-identification in place but achieves no point-identification. When the discriminating question coincides with the disputed margin, point-identification requires a further restriction A' , fixed outside the dispute, on which only one mechanism is consistent: an over-identifying anchor of the kind Part III supplies through the borrower's tier. A *machloket* of the identifying kind is a pair of positions $\{(M_1, a_1), (M_2, a_2)\}$ recorded against the shared rule R . Observing the dispute and the subsidiary rulings identifies the mechanism: the law that rules $A = a_1$ is generated by M_1 . In the binary branch the mechanisms are mutually exclusive and the recorded ruling selects one; in the nested augmentation $M_2 = M_1 \oplus C$ for an added ground C , the mechanisms stand in an inclusion, and the ruling identifies whether C is present. Resolution of the dispute, where it occurs, is the law's adoption of one (M_i, a_i) , and so records the mechanism selected as well as the mechanisms available.