

The Vulnerability Premium: Why Fairness Norms and Consumer Law Both Miss State-Based Pricing

Shay Tsaban

shay.tsaban@gmail.com

Peres Academic Center

Research Article

Keywords: price fairness, personalized pricing, consumer vulnerability, fairness theory, dual entitlement, consumer protection

Posted Date: September 9th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10738984/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

European consumer regulation identifies vulnerable consumers by stable characteristics: infirmity, age, credulity, disability, or economic situation. Behavioral research locates much impaired decision-making elsewhere, in transient states such as acute emotion, fatigue, time pressure, and scarcity, which vary within a person over hours rather than across persons over years. Digital markets can infer such states from behavioral traces and condition prices on them. This article names the resulting increment the vulnerability premium: the part of an offered price attributable to an inference about the consumer's transient state rather than about the consumer's stable type. The premium escapes both systems that ordinarily discipline pricing. Fairness norms fail first. Research from dual entitlement to fairness theory shows that fairness judgment runs on reference transactions and counterfactual comparisons, and the premium is constructed so that neither a reference nor a counterfactual is available to the person paying it. The moralized market reaction that constrains visible price discrimination therefore cannot arise. Consumer law fails second, for five separate doctrinal reasons: a group requirement, a closed trait list, a distortion test that the practice's universality defeats, a significant-harm threshold, and a disclosure duty discharged at the moment of lowest capacity. Three design responses would restore the conditions of fairness judgment: state-blind pricing inputs, a right to be re-quoted, and evidentiary burden shifting. The argument extends justice research to a pricing practice built to be imperceptible, and it specifies what would falsify it.

Introduction

A consumer who shops at two in the morning after a long shift is not a different person from the one who shops on Saturday after breakfast. The law treats them as the same consumer, because the law identifies consumers by what they are rather than by how they are. A pricing system does not have to make that assumption. It observes the hour, the scroll velocity, the hesitation before the checkout button, and the three abandoned carts earlier in the evening. It can then act on what those traces imply about the state the person is in right now.

Two systems ordinarily stand between a consumer and an exploitative price. The first is informal. Markets are moralized. Buyers judge prices against reference transactions, treat unjustified increases as violations of an entitlement, and withdraw custom from firms that violate it, which is why fairness operates as a constraint on profit seeking even where law is silent (Kahneman et al., 1986). The second is formal. European consumer law contains a developed apparatus for protecting vulnerable consumers against practices that exploit them. This article is about a pricing practice that passes through both systems, and about why it passes through each for a structurally similar reason.

Both systems identify the person at risk by what that person durably is. The Unfair Commercial Practices Directive protects consumers who belong to a clearly identifiable group and who are particularly vulnerable because of "their mental or physical infirmity, age or credulity". The Artificial Intelligence Act prohibits the exploitation of vulnerabilities due to "age, disability or a specific social or economic

situation". Fairness research, for its part, has anchored the judgment of prices in comparisons across persons and across transactions: the reference transaction of dual entitlement, the comparison other of equity theory, the referent of relative deprivation. In each case the machinery presupposes a stable someone, or a stable something, to compare against.

Behavioral research points somewhere else. A large body of work locates a substantial share of impaired decision-making in states rather than in types. Visceral influences shape choice in ways a person cannot fully anticipate beforehand or recall afterwards (Loewenstein, 1996). Emotions reshape risk perception and time preference (Lerner et al., 2015; Loewenstein et al., 2001). Scarcity has cognitive consequences of its own (Shah et al., 2012), as do fatigue and time of day (Killgore, 2010; Wieth & Zacks, 2011). These are not characteristics of persons. They are characteristics of persons at times. The same individual moves in and out of them repeatedly, often within a single day, and the movement is not observable to the individual with any precision.

The distinction between what a person is and how a person is has a settled place in vulnerability theory. Fineman (2008) argued that vulnerability is universal and constant, inherent in the human condition rather than confined to identifiable groups, and used that claim to press for a more responsive state. Rogers et al. (2012) distinguished dispositional from occurrent vulnerability, the standing susceptibility from the realized episode. Cole (2016) added that invoking vulnerability is best understood as advancing a claim about injustice rather than as describing a group's weakness. The argument developed here takes that last observation literally. If vulnerability claims are injustice claims, then a market practice that converts occurrent vulnerability into revenue belongs on the agenda of justice research, not only on the agenda of lawyers. Vulnerability on the present account is universal but not constant. It arrives and departs, several times a week, and its arrival is legible to anyone holding the right data. What can be detected can be priced.

The commercial significance of that mismatch has changed. When a trader could observe only the aggregate characteristics of a customer segment, a trait model of vulnerability was a reasonable approximation of what traders could exploit and therefore of what norms and law needed to police. Personalization removes the approximation. A system that records the timing, sequence, and hesitation structure of a session holds exactly the material that reveals transient states. Inference of behavior and affect from passively collected traces is an established research program, not a speculation (Insel, 2017). Commercial emotion-inference systems are established enough to have attracted their own regulatory literature (McStay & Rosner, 2021), and the continuous adjustment of an environment to the individual has been analyzed as a distinct mode of regulation by design (Yeung, 2017). Once the firm's behavioral model of the consumer is finer than the protector's, protection built on the protector's model inherits its coarseness, a general problem for behavioral policy in the age of tailored intervention (Mills et al., 2023).

This article names the resulting price increment the **vulnerability premium**: the part of an offered price that is attributable to an inference about the consumer's transient state rather than to an inference about

the consumer's stable preferences, costs of service, or willingness to pay. The thesis can be stated in one sentence: state-based pricing is a fairness violation engineered to be imperceptible, escaping fairness norms because they presuppose a perceivable violation, and escaping consumer law because it presupposes an identifiable victim.

The article makes three contributions. It defines the premium as a distinct object, with defining attributes, a decision rule separating it from its neighbors, and stated conditions under which the construct would be redundant. It explains why the practice escapes informal fairness enforcement, drawing on dual entitlement, equity theory, fairness theory, and uncertainty management. It explains separately why the practice escapes European consumer law, through five doctrinal gates that fail for five different reasons.

It then derives three regulatory instruments that are best read as restoring the informational conditions under which fairness judgment and legal protection can operate at all. The argument proceeds in that order, and it closes with the confounders and findings that would weaken it, because a construct offered to an empirical field should say what would count against it.

Vulnerability as Trait and Vulnerability as State

The Trait Model in Protective Regulation

European consumer law assesses commercial practices from the perspective of the average consumer. The Court of Justice fixed the benchmark in *Gut Springenheide* as a consumer "reasonably well informed and reasonably observant and circumspect", and the Unfair Commercial Practices Directive later codified it. The benchmark is a disposition, and dispositions are stable by construction. A consumer who is reasonably observant is reasonably observant on Tuesday and on Saturday, at nine in the morning and at midnight. Nothing in the benchmark contemplates that the same person satisfies it at one hour and fails it at another.

The Directive does contain a corrective. Article 5(3) sets a special benchmark for practices likely to materially distort the economic behavior "only of a clearly identifiable group of consumers who are particularly vulnerable to the practice or the underlying product because of their mental or physical infirmity, age or credulity". Where the trader could reasonably be expected to foresee the distortion, the practice is assessed from the perspective of the average member of that group. The provision has been read as the Directive's principal concession to the behavioral critique of the average consumer standard (Duivenvoorde, 2015), and it has carried much of the weight in subsequent arguments about protecting consumers in digital markets (Riefa, 2022). Trzaskowski (2011) observed early that the Directive absorbed behavioral economics selectively, importing the finding that consumers deviate from the rational benchmark while leaving the deeper architecture of the average consumer intact.

Read closely, Article 5(3) imposes four conditions, and each presupposes that vulnerability is a property of persons rather than of moments. The consumers must form a group. The group must be clearly identifiable. Membership must follow from infirmity, age, or credulity, which is a closed list of durable characteristics. And the trader must be able to foresee the distortion, which is easiest to establish when the characteristic is visible in advance. The provision is an instrument for protecting categories of people, and it is well designed for that purpose.

The trait model is not confined to the Directive's text. It shapes how vulnerability is operationalized in sectoral practice, where the standard approach is to identify consumers who possess a durable characteristic and then attach obligations to dealings with them. Cartwright's (2015) taxonomy of vulnerability in financial services distinguishes informational, pressure, supply, redress, impact, and market vulnerability, and treats each as an attribute the consumer brings to the transaction. The taxonomy contains no category for a consumer who possesses none of these attributes on Tuesday morning and several of them at midnight. The practices it was built to describe did not require one.

The doctrine that comes closest to the moment is the prohibition of aggressive practices. Article 2(j) of the Directive defines undue influence as "exploiting a position of power in relation to the consumer so as to apply pressure", in a way that significantly limits the ability to make an informed decision. The Court's two leading judgments mark the provision's reach and its limit. In *Wind Tre*, the sale of SIM cards with services preloaded and preactivated, without adequate prior information, was held to be an aggressive practice, because the trader had imposed a default on everyone. In *Orange Polska*, by contrast, requiring the consumer to sign in the physical presence of a waiting courier was held not to be aggressive in itself, absent additional conduct that applied pressure. The pattern matters here. The doctrine catches structures a trader builds and pressure a trader applies. A trader who prices a consumer's own passing state builds nothing, applies nothing, and waits. There is no conduct for the aggression doctrine to attach to, only an inference and a number.

The Artificial Intelligence Act repeats the trait structure in a newer setting. Article 5(1)(b) prohibits AI systems that exploit vulnerabilities "due to their age, disability or a specific social or economic situation", where the exploitation distorts behavior and causes or is likely to cause significant harm. Age and disability are stable. A specific social or economic situation is the broadest of the three and reaches, for example, poverty, but it is a situation a person is in across sessions rather than a state a person is in tonight.

One feature of the drafting cuts the other way and should be stated plainly, because it narrows the argument made here. Article 5(1)(a), which prohibits subliminal and purposefully manipulative techniques, applies where such techniques have the objective "or the effect" of materially distorting behavior. The disjunction means the Act does not require proof of intent to manipulate. What survives is narrower: the technique itself must still be characterized as manipulative, deceptive, or subliminal, and the harm must clear a significance threshold. A price a few percentage points higher than it would

otherwise have been is not obviously the product of a manipulative technique, and it is not obviously a significant harm.

The most direct European response to personalized pricing is an information duty, and its premise is the most revealing text in this area. Directive (EU) 2019/2161 requires the trader to inform the consumer, where applicable, that the price was personalized on the basis of automated decision-making. The recital that justifies the duty explains that profiling allows traders “to assess the consumer’s purchasing power”. The premise is explicit. Personalization is understood as an inference about a stable economic attribute of the person, which matches how the economics of price discrimination had framed the problem (Acquisti et al., 2016) and how early legal analysis framed it (Zuiderveen Borgesius & Poort, 2017). It is a type model, and it produces a type-model remedy. If the trader has estimated how much money a consumer has, telling the consumer that an estimate was made is a sensible response, because consumers know how much money they have and can compare. The recital also carves out dynamic real-time pricing that responds to market demand, which creates an obvious characterization problem for any practice that varies prices rapidly on signals correlated with time and context.

What Behavioral Research Locates in States

The behavioral claim relevant here is old and well established. Loewenstein (1996) argued that visceral factors such as hunger, fatigue, pain, craving, and strong emotion exert an influence on behavior that people systematically underestimate when they are not experiencing them, and misremember once the state has passed. The asymmetry is the point. A person in a hot state does not have accurate access to how the state is shaping their valuation, and a person in a cold state cannot easily predict it.

Work on emotion and decision-making has specified the mechanisms in some detail. Lerner et al. (2015) synthesize evidence that incidental emotions, meaning emotions unrelated to the decision at hand, carry over into judgment and alter risk perception, time preference, and depth of processing. Loewenstein et al. (2001) had earlier set out the case that anticipatory feelings drive risky choice through a channel partly independent of cognitive evaluation of outcomes.

Sleep loss degrades exactly the functions consumer decisions rely on. Killgore (2010) reviews evidence that sleep deprivation impairs attention, working memory, and higher-order executive function, with larger effects for tasks requiring the inhibition of a prepotent response. Frenda et al. (2016) show that a night of total sleep deprivation substantially raises the probability that a person will sign a statement they believe to be false, a striking demonstration that a transient state can alter consent-like behavior rather than only performance.

Time of day matters in ways that are not simply monotonic. Wieth and Zacks (2011) found that performance on some problem types is better at a person’s non-optimal time of day, a useful corrective against assuming that late-night consumers are uniformly worse at everything. The claim needed here is

weaker and more secure: decision quality varies systematically within a person across the day, and the variation is legible in behavioral traces.

Scarcity operates similarly. Shah et al. (2012) show that having too little of a resource induces a focus on the immediate problem that borrows attention from other domains, producing choices that look like traits of the poor but are consequences of the situation.

The Evidential Caution

A regulatory argument that leans on behavioral findings should say which of those findings are contested. Two cautions belong here. The first concerns the best-known demonstration of state-dependent decision-making, the finding of Danziger et al. (2011) that judicial parole decisions varied with the timing of food breaks. The result was challenged shortly after publication on the ground that case ordering was not random (Weinshall-Margel & Shapard, 2011). The exchange is unresolved enough that the finding should not carry weight in a regulatory argument, and it carries none here. The second concerns the wider literature on self-control as a depletable resource, which has not replicated reliably. The argument in this article does not require depletion. It requires only that decision-relevant capacities vary within a person over short horizons, and that the variation is correlated with observable behavior. Fatigue, acute emotion, and scarcity each supply that, and none of them depends on the depletion account being correct.

The Vulnerability Premium

Definition and Defining Attributes

The vulnerability premium is the component of an offered price that is attributable to an inference about the consumer's transient state, rather than to an inference about the consumer's stable preferences, the cost of serving them, or general market conditions at that moment.

Four attributes define it, and all four must be present.

The priced attribute is transient. It is a property of the consumer at a time, and the same consumer moves out of it without doing anything about it.

It is also endogenous to the transaction environment. The traces that reveal the state are produced by the act of shopping. A consumer cannot browse without generating timing, sequence, and hesitation data, and cannot suppress the signal without abandoning the market.

The consumer cannot observe the premium. Prices are offered, not decomposed. A consumer sees one number and has no access to the number that would have been offered in a different state. The

counterfactual is not displayed, and unlike a geographic or device-based difference it cannot be recovered by asking a friend to check.

The state impairs the capacity that any remedy assumes. Whatever the consumer is supposed to do in response, whether reading a disclosure, comparing offers, or deferring the purchase, is made harder by the very state that triggered the premium.

The third and fourth attributes are what distinguish this from an ordinary pricing grievance and are the reason the construct earns a name. In the standard account of an information remedy, the consumer is uninformed and the disclosure informs them. Here the disclosure would arrive at the moment of lowest capacity to use it, and it would arrive because capacity is low.

The Nearest Neighbor: Digital Vulnerability

The account against which this construct must earn its place is digital vulnerability. Helberger et al. (2022) argue that vulnerability in digital markets is architectural and universal rather than a property of particular groups, because choice architectures that are data-driven, dynamically adjustable, and personalizable can identify or induce susceptibility in anyone. The fullest statement is the report by Helberger, Lynskey, Micklitz, Rott, Sax, and Strycharz (2021), which grounds the concept in the vulnerability taxonomy of Rogers et al. (2012): sources of vulnerability may be inherent or situational, and its states may be dispositional or occurrent. The report concludes that in data-driven markets every consumer is dispositionally vulnerable, and that the vulnerability is relational, produced by the ongoing relationship between consumer and architecture. Consumer policy should therefore stop sorting people into vulnerable and non-vulnerable groups and start scrutinizing the practices that render people vulnerable. Its central reform proposals are structural: treating exploitative data practices and structural asymmetry as unfair, and reversing the burden of proof.

The trait-state contrast, taken alone, is therefore not new. Vulnerability theory has distinguished the dispositional from the occurrent for over a decade, and the digital vulnerability literature has imported that distinction into consumer policy with considerable sophistication. The present construct is not a rival to that framework. It is a specialization of one cell of it, the cell the framework identifies and then passes over: the occurrent episode, in the moment of its occurrence, considered as an input to a price.

Three things follow from that specialization that the structural account does not supply. The first is an object. Digital vulnerability is a condition of persons in an architecture, and conditions cannot be counted. The premium is a component of a price, which makes the exploitation quantifiable in principle and gives an authority something to look for. The second is a set of identification conditions, developed below, which state what an observer would have to see before concluding that a price contains a premium, and which turn on within-person rather than between-person comparisons. The third is a mechanism of impunity. The structural account explains why exploitation is possible: the architecture holds the data and adapts. It does not explain why this particular form goes unpunished by the informal fairness enforcement that visibly disciplines other pricing practices. That explanation, developed in the

next section, is the construct's main analytical payoff, and it is a justice-theoretic explanation rather than a legal one.

The difference has practical consequences. A rule addressed to structural asymmetry does not by itself prohibit conditioning a price on tonight's fatigue, and a reversed burden of proof helps an authority that already suspects, not a consumer who cannot perceive. Establishing that a market is architecturally asymmetric establishes that any given price might contain a premium. It does not establish that any particular price does, and the distance between those two claims is where enforcement lives.

Manipulation, Willingness to Pay, and Market State

Three further neighbors need separating.

The first is market manipulation and its digital successors. Hanson and Kysar (1999) established the claim that firms exploit predictable cognitive limitations, and Calo (2014) updated it for a marketplace in which the firm designs the environment and observes the response. Sunstein (2016) supplies the definitional apparatus, treating a practice as manipulative when it fails to engage the target's capacity for reflection, and Susser et al. (2019) sharpen it further: manipulation is hidden influence on decision-making. Adjacent work applies the frame to unfair commercial practices doctrine (Hacker, 2023), to the erosion of transactional autonomy online (Mik, 2016), and to choice architecture and dark patterns specifically (Brenncke, 2024; Luguri & Strahilevitz, 2021; Mathur et al., 2019).

All of these accounts share a structural feature: the trader acts on the consumer's decision process. A dark pattern alters the choice architecture in order to change what the consumer does. The vulnerability premium leaves the architecture alone and changes the price. A consumer facing a premium may be making a perfectly clear-eyed decision on an interface with no manipulative features at all, and paying more for it. The influence is hidden, as in Susser et al.'s definition, but nothing is influenced except the terms. That is why the manipulation frame, otherwise the closest home, does not fit: nothing about the decision has been distorted. Only its cost has changed. Where a system both distorts the interface and prices the state, the two are present together and should be analyzed separately.

The second is ordinary personalized pricing, understood as price discrimination on inferred willingness to pay (Bar-Gill & Sunstein, 2025; Seele et al., 2021; Zuiderveen Borgesius & Poort, 2017). The distinction is not the technology, which is identical, but the object of the inference. Willingness to pay is a preference. A transient state is a capacity. Charging a wealthy consumer more because they will pay more raises distributive questions that are at least intelligible in welfare terms, and that firms can defend in those terms. Charging a tired consumer more because tiredness raises the reservation price of ending the search is a different transaction, because the elevated reservation price is an artifact of the state and does not survive it. Philosophical accounts of exploitation describe exactly this structure. An advantage extracted through a consensual and even mutually beneficial exchange can still be wrongful when it trades on the other party's vulnerability (Goodin, 1985; Wertheimer, 1996). More recent accounts locate

the wrong in the use of another's vulnerability as such, whether or not the victim ends up worse off (Vrousalis, 2013).

Behavioral analyses of contract design have shown how firms price against consumers' systematic misperceptions (Bar-Gill, 2012). The premium extends that logic from stable misperception to momentary incapacity. Related work has asked whether personalized marketing communication exploits vulnerability (Strycharz & Duivenvoorde, 2021), and how applications convert an intimate behavioral stream into commercial advantage (Sax et al., 2018). Both address what is communicated to the consumer rather than what is charged.

The third is dynamic pricing responding to demand. The difference is the unit at which the signal is measured. Dynamic pricing conditions on market state, and everyone facing the market at that moment sees the same number. The premium conditions on personal state, and two consumers in the same market at the same moment see different numbers for reasons that have nothing to do with the market.

This yields a decision rule. An observed price contains a vulnerability premium when four conditions hold together. The price varies across consumers facing identical market conditions and identical costs of service. The variation is predicted by trace-derived indicators of transient state rather than by indicators of stable type. The same consumer receives a different price in a different state. And the differential does not survive the state. The third and fourth conditions are what make the claim falsifiable, since both are within-person comparisons.

The construct should be treated as redundant if any of three results holds. If price variation attributable to state indicators disappears once stable type is controlled for, the phenomenon is willingness-to-pay discrimination measured with noisy proxies. If within-person price differentials track only time since last purchase or other commercial-relationship variables rather than state, the phenomenon is ordinary customer-lifecycle pricing. And if consumers who are told the counterfactual price behave no differently from those who are not, the informational asymmetry the construct rests on is not doing the work claimed for it.

Working the Metaphor Against Its Source

The word premium is borrowed from insurance and finance, where it names a payment made in exchange for the transfer of risk, or the excess of a price over a benchmark value. Borrowed vocabulary imports default inferences, and the borrowing should be tested rather than assumed.

Two features carry over. A premium is a differential measured against a benchmark, and the construct here is likewise defined by reference to a counterfactual neutral-state price. A premium is also priced by a party with better information about the relevant risk than the payer has, which is precisely the informational structure at issue.

Three features do not carry over, and naming them is what converts the metaphor into a model. An insurance premium is consideration for something the payer receives, namely the assumption of risk by the insurer, whereas the vulnerability premium purchases nothing. An insurance premium is disclosed and contractually agreed, whereas this one is embedded in a single undecomposed number. And an insurance premium is calculated from a stable risk classification that the insured can in principle contest, whereas this one is calculated from a state that will have passed before any contest could be brought. The third disanalogy is the most consequential, because it explains why remedies modeled on the correction of misclassification, such as rights of access and rectification, do not fit.

What Is Identified and What Is Not

The level of the vulnerability premium in any individual transaction is not identified, and no method proposed here identifies it. A researcher or regulator observing one price does not observe the counterfactual, and the firm's own estimate of the state is not a fact about the consumer but a model output.

What is identified is more limited and still sufficient. The sign is identified where within-person variation can be observed, since the direction of the differential between a consumer's own prices in different states can be established without knowing the neutral baseline. The ordering across consumers is identified under an assumption of a common pricing function. Change within a consumer over time is identified directly. And bounds are identified, in the sense that the observed dispersion of prices offered to the same consumer under identical market conditions places a lower bound on the magnitude of state-conditioned variation.

The justice-evaluation tradition makes the stakes of this limitation exact. Jasso's (1980) justice evaluation is a function of two quantities, the actual reward and the just reward, and it is undefined when either input is missing. Justice research has spent four decades refining how people form the second quantity (Moya & Adriaans, 2026). The premium attacks the first. A consumer who never learns the actual differential cannot compute an evaluation at all, however refined their standard of the just price. Any single figure offered as the value of the premium for a market or a population should therefore be treated as an assumption presented as an estimate. This article does not license one, and a regulatory design that depends on producing one is for that reason a weak design. The proposals below are constructed to avoid the dependence.

Why Fairness Norms Cannot Discipline It

Fairness as a Constraint on Profit Seeking

The informal system that polices pricing has been documented for four decades. Kahneman et al. (1986) showed that community standards assign firms an obligation to respect a reference transaction: raising the price of snow shovels after a blizzard is judged unfair by large majorities, and firms anticipate the judgment and hold prices below what the market would bear. The mechanism is dual entitlement.

Customers are entitled to the terms of the reference transaction, firms to their reference profit, and a price increase is acceptable when it defends the reference profit against a cost shock but unfair when it merely exploits a shift in the buyer's circumstances. The enforcement is informal and effective: perceived unfairness produces complaint, exit, and reputational damage, and the anticipation of that reaction constrains pricing even where no rule applies.

The price-fairness literature has specified the judgment's inputs. Consumers evaluate a price against reference points drawn from their own past transactions, from other buyers' prices, and from inferences about the seller's costs and motives (Bolton et al., 2003). Judgments of unfairness are comparative by construction: a price is unfair relative to a comparison, and the comparison's similarity determines the judgment's force (Xia et al., 2004). Dynamic and differential pricing is judged more harshly precisely when the variation tracks the buyer rather than the product, and fairness perceptions deteriorate as the same good is seen to carry different prices for different people at the same moment (Haws & Bearden, 2006). The acceptability of a pricing policy depends on the justification available for it, a regularity documented for public pricing as well (Schuitema et al., 2011). The pattern across this literature is consistent. Fairness enforcement is triggered by a perceived deviation from a perceivable benchmark.

The literature on personalized pricing confirms that the trigger works when it is pulled. Attitudes toward online price discrimination are strongly negative once the practice is understood (Poort & Zuiderveen Borgesius, 2019). Willingness to accept personalization is low once it is made salient (Leszczynska & Baltag, 2024; Li & Nill, 2020). van Boom et al. (2020) showed experimentally that a mandated disclosure of discriminatory pricing shifts intention to purchase, and that the shift depends on how the disclosure is framed. Consumers who are told that a price was set for them react, and the reaction is strong enough that firms have reason to fear it. The moralized market is real.

Fairness Judgment Runs on References and Counterfactuals

Justice research is precise about what a fairness judgment consumes. Equity theory requires a comparison other whose ratio of outcomes to inputs can be set against one's own (Adams, 1965). Relative deprivation requires a referent: resentment arises from comparing one's situation to that of another person, group, or time, and appraising the difference as undeserved (Smith et al., 2012). Fairness theory locates the judgment in counterfactual thought: an event is unfair when the person can construct an alternative that would have been better, that the agent could have brought about, and that the agent should have brought about (Folger & Cropanzano, 1998). Counterfactual availability is not a nicety of the theory but its engine, and experimental work has confirmed that manipulating the ease of counterfactual thought changes fairness perceptions (Nicklin, 2013).

The same research program has mapped what happens when the inputs are missing. When people lack information about others' outcomes, they do not suspend judgment. They substitute, drawing on procedural information to construct a fairness judgment (van den Bos et al., 1997). Under uncertainty, fairness judgments serve as heuristics, assembled from whatever information arrives first and is

available at all (Lind & van den Bos, 2002). Even outcome fairness can stand in for procedural fairness when procedural information is inaccessible (Daly & Tripp, 1996). The substitution machinery is resourceful, and it needs at least one input to run on.

State-based pricing removes every input at once. There is no reference transaction, because a personalized price is not posted anywhere and each buyer's history is their own. There is no comparison other, because two consumers in the same market at the same moment see different numbers and neither sees the other's. There is no counterfactual, because the neutral-state price is never displayed and, unlike a geographic difference, cannot be recovered by asking a friend to check. And the procedural fallback is blocked twice over: the pricing procedure is a trade secret, and the consumer does not know that a procedure was applied to them at all. The substitution machinery documented by uncertainty management research cannot run, not because uncertainty is high but because the event itself is not represented. Nothing presents itself for judgment.

Engineered Imperceptibility

The conclusion is not that consumers judge the vulnerability premium to be fair. It is that the judgment never occurs, and the distinction carries the argument. Dual entitlement constrains firms through anticipated reaction, and the reaction requires a perceived violation. A firm that prices tonight's fatigue does not risk the snow-shovel reaction, because the buyer has no benchmark against which tonight's number reads as a violation of anything. The disclosure experiments make the dependence visible from the other side: when the practice is disclosed, reactions follow, and the size of the reaction depends on the framing of the disclosure (van Boom et al., 2020). A practice whose profitability survives only in the absence of the fairness reaction is a practice with a commercial interest in imperceptibility, and state-based pricing has that interest built into its defining attributes. The third attribute removes the benchmark. The fourth degrades the capacity to seek one.

Justice research has begun to insist that fairness is enforced socially rather than computed privately. Lind (2023) argues that the field's cognitive turn understated the social dimension of justice, the shared seeing and collective reaction through which fairness standards actually bind. Lerner (2023) distinguishes the expression of social norms from the justice motive proper, and both presuppose that the triggering event is available to be seen. Tyler (2023) traces how justice theory developed in organizational settings, where allocations and procedures are at least in principle observable and contestable by the people they govern.

The vulnerability premium is a boundary case for all three accounts. It is an allocation decision with no audience: not concealed in the manner of fraud, which misrepresents a fact that could be checked, but structured so that the fact never takes a form a person could check. A field whose instruments measure perceived injustice meets, in this practice, a phenomenon in which injustice and its perception come apart by design. That is what makes the premium a research object for justice scholarship and not only a regulatory nuisance: it marks the difference between studying fairness reactions and studying the conditions under which fairness reactions are possible.

The point also explains why the two neighboring literatures each stop short. The manipulation literature attends to hidden influence on decisions, and here the decision is untouched. The price-discrimination literature attends to willingness to pay, and here the priced quantity is a capacity, not a preference. Between them lies a practice that neither frame names: the conversion of a person's worst hour into a margin, at scale, without a single distorted choice or a single perceivable comparison.

Why Law Cannot Substitute

If informal enforcement cannot see the practice, formal protection has to carry the weight alone. It cannot, and the claim here is specific. The phenomenon passes through five distinct legal gates, each for a different reason, which is why patching any one of them would not be sufficient.

The Group Requirement

Article 5(3) of the Unfair Commercial Practices Directive protects a clearly identifiable group. A transient state does not constitute a group in any sense the provision can use. The set of consumers who are fatigued at a given moment has no stable membership, no shared characteristic that persists, and no boundary that a trader or an enforcement authority could draw in advance. Any given consumer is inside it several times a week and outside it the rest of the time.

The Closed Trait List

Even if the group requirement were relaxed, the qualifying characteristics under the Directive are mental or physical infirmity, age, and credulity, and under Article 5(1)(b) of the Artificial Intelligence Act they are age, disability, and a specific social or economic situation. Ordinary tiredness is not an infirmity, and a person who is upset at eleven at night is not thereby disabled. Treating transient states as infirmities would medicalize ordinary variation in a way that would be both descriptively wrong and, applied consistently, absurd.

The Word Only

The most striking feature of Article 5(3) is easy to miss. It applies to practices likely to materially distort the economic behavior “only” of the protected group. The condition is a limiting device with a defensible purpose, since it prevents a general commercial practice from being assessed against the standard of its most susceptible audience.

Applied to state-based pricing it produces a perverse result. Because everyone is sometimes tired, sometimes rushed, and sometimes upset, a practice that prices those states distorts the behavior of the whole population, some of the time. It therefore fails the only-requirement by being too widely applicable, and the protection weakens precisely as the practice becomes more general. Universality of exposure defeats the protection rather than triggering it. Here the digital vulnerability literature and this article converge from different directions. If every consumer is dispositionally vulnerable, as Helberger et al. (2021) conclude, then a provision reserved for practices that distort only a delimited group's behavior is structurally incapable of protecting anyone from practices aimed at the vulnerability everyone shares.

This is the clearest instance in the analysis of a rule whose behavioral premise, that vulnerability is concentrated in a minority, is written into its operative wording.

The Significant Harm Threshold

Both limbs of Article 5(1) of the Artificial Intelligence Act require harm that is significant, or reasonably likely to be significant. A modest price differential on a single transaction will rarely qualify, and the aggregation problem is severe. Harm that is small per transaction, repeated across many transactions and many consumers, and invisible in each instance, is the standard case in which a significance threshold under-enforces.

Disclosure at the Moment of Lowest Capacity

The personalized pricing duty under the Consumer Rights Directive is an information remedy, and information remedies have a well-documented record of failure in consumer markets (Ben-Shahar & Schneider, 2014). The failure mode here is specific rather than general. A disclosure works if the recipient can act on it. Acting on this one requires noticing the notice, understanding that the price reflects an inference about oneself, forming a view about whether that inference is being made now, and then deferring or comparing rather than buying. Every element of that sequence draws on capacities that the triggering state degrades. The remedy is administered at the moment of lowest capacity to use it, and the timing is not incidental. It is causally downstream of the same state that produced the premium.

There is a further difficulty specific to the wording. The duty requires the trader to state that the price was personalized, not on what basis, and the recital's framing in terms of purchasing power invites the consumer to interpret the notice as being about their means. A consumer who reads the notice and concludes that the trader has estimated their income has read it exactly as the legislator described it, and has drawn the wrong conclusion about what is happening. The disclosure-framing results discussed above bear directly on this point: reactions to a personalization notice depend on what the notice leads the consumer to believe was personalized (van Boom et al., 2020). An evaluation of the new pricing disclosures accordingly found modest effects on purchase intentions (Barrafrem & Tinghög, 2026). None of that evidence concerns state-based pricing specifically, which is itself part of the argument, since the phenomenon has not been separated out for study.

The Sum of the Gates

The Commission's own review effort has documented the pattern from inside. The fitness check of EU consumer law on digital fairness concluded that the existing directives remain necessary but leave gaps around dark patterns, addictive design, and unfair personalization practices (European Commission, 2024). The average consumer benchmark and the group-based vulnerability provision remain the frame within which those gaps are described. The BEUC report goes further, proposing that structural digital asymmetry itself be treated as unfair and that the burden of proof be reversed so that traders must show their data practices are not exploitative (Helberger et al., 2021).

Those proposals would repair the asymmetry between firms and enforcement authorities, and the burden-shifting instrument developed below stands in their lineage. What they do not reach is the asymmetry between the firm and the consumer at the moment of purchase. A consumer cannot invoke a reversed burden against a number that never announces itself as a differential. The doctrinal gates and the perceptual failure are therefore not two problems but one problem seen from two sides. Every gate assumes a protected class, a caught practice, or an informed consumer, and the premium is constructed so that there is no class, no conduct, and no information.

Restoring the Conditions of Fairness

Regulatory design can be read as a form of justice engineering, the deliberate construction of market institutions whose outcomes people can recognize as fair (Lotz, 2015). Read that way, the design task here is not to forbid a technology but to restore the informational conditions under which fairness judgment and legal protection can operate. Three instruments follow, ordered from prevention to proof.

State-Blind Pricing Inputs

The simplest instrument regulates the input rather than the outcome. A rule prohibiting the use of specified categories of trace data as inputs to price determination would not require anyone to measure a premium, prove significant harm, or identify a group.

The categories would need to be defined by function rather than by name, since any list of variables invites substitution. A workable formulation prohibits conditioning price on inferences about a consumer's momentary affective, attentional, or physiological condition. It would then treat a set of proxies as presumptively caught: local time of day at the consumer's location, session duration, input latency and correction rates, and repeated abandonment within a session. Traders would remain free to use these variables for fraud detection, security, and service quality, which is where their legitimate value lies, and the restriction would attach to price determination only.

The design has an evident cost. Some of the prohibited variables carry genuine cost or demand information, and forbidding their use in pricing forgoes efficiency. The trade is defensible where the variable's predictive power for price runs mainly through the state channel. Stated honestly, the proposal accepts a loss of pricing precision in exchange for removing a class of transaction that neither market norms nor law can discipline.

A Right to Be Re-Quoted

The second instrument works with the transient nature of the state rather than against it. If the state passes, a remedy that lets the consumer transact after it has passed restores the position without anyone having to measure anything.

A right to be re-quoted would entitle a consumer to obtain, on request and without charge, a fresh price for the same product at a specified later time, with the trader bound to honor the lower of the two for a

short window. The instrument is a close analogue of the withdrawal right, which already embodies the insight that the moment of purchase is often not the moment of best judgment. It inherits that instrument's principal advantage: it requires no inquiry into the trader's reasons.

In the terms of this article, the right manufactures the missing counterfactual. A consumer holding two quotes for the same product, one from last night and one from Saturday morning, holds exactly the within-person comparison that fairness judgment runs on and that the practice is structured to withhold. The remedy does not tell the consumer whether the difference was fair. It restores the conditions under which that question can be asked, by the person entitled to ask it. Fairness theory would describe this as making counterfactual thought possible again; the justice-evaluation tradition would describe it as supplying the actual reward that the evaluation function requires.

Two objections deserve answers. The first is that traders would inflate later quotes. Binding the trader to the lower of the two prices removes the incentive. The second is that the right would be used rarely, which is the fate of most consumer rights. That is likely true at the level of individual use, and it is not the main channel of effect. A trader who must honor a neutral-time price on request faces a reduced return to state-based pricing in the first place, because the differential becomes recoverable by any consumer who asks.

Evidentiary Burden Shifting

The third instrument addresses proof. Because the counterfactual price is unobservable to the consumer and to the regulator, an enforcement regime that requires the claimant to establish the premium will not function.

A workable rule shifts the burden once a threshold showing is made. If an authority demonstrates that a trader's prices for identical products under identical market conditions vary systematically with state-correlated indicators, the trader is required to show that the variation is explained by costs, by market conditions, or by stable consumer characteristics. Traders hold the pricing model, the feature list, and the experimental record, so the information is available to the party asked to produce it. This mirrors the structure familiar from discrimination law, where a *prima facie* showing of differential treatment shifts the explanatory burden to the party with access to the reasons. It extends the burden-reversal logic already proposed for data exploitation generally (Helberger et al., 2021) to the specific evidentiary problem of state-based pricing.

What the Digital Fairness Act Would Have to Decide

The European Commission is preparing a Digital Fairness Act, with a proposal expected in late 2026, and its announced remit includes unfair personalization practices, particularly where consumer vulnerabilities are exploited for commercial purposes. Three drafting choices will determine whether the instrument reaches the premium or reproduces the trait model. The first is whether vulnerability is defined by enumerated characteristics or by function. A list will inherit the closed-list problem no matter how long it grows, because the relevant class is defined by transience rather than by content. A

functional definition, framed in terms of circumstances in which a consumer's capacity to make an informed decision is temporarily reduced, does not have that defect.

The second is whether the instrument retains a requirement that the practice affect a delimited group, which would reproduce the result that the broadest practices attract the weakest scrutiny. The third is whether the instrument relies on disclosure, which would administer the remedy at the moment of lowest capacity. None of these is a technical drafting point. Each follows from whether the drafter's model of the consumer is a type or a person in a state.

Limits, Confounders, and What Would Falsify the Argument

Several things could be true that would weaken this argument, and they should be stated rather than left for a reader to supply.

The most important is that the practice may be less prevalent than the capability suggests. This article establishes that the inference is feasible, that the incentive exists, and that neither fairness norms nor law would catch the practice. It does not establish how much of it occurs. Capability arguments have a poor record of predicting commercial behavior, and a firm may decline to price on state for reputational reasons, because the gains are small relative to the risk, or because its pricing model is too coarse to support it. The honest characterization is that this is an argument about a gap that is open, not a report of a practice that has been measured.

Attribution is genuinely difficult. Prices vary at night for reasons that have nothing to do with the consumer, including thinner inventory, different competitor behavior, and demand composition. A late-night price differential is therefore not evidence of a premium, and any empirical program has to separate the market channel from the personal channel. Within-person designs, in which the same consumer is observed under identical market conditions in different states, are the only designs that do this cleanly, and they are difficult to run without the cooperation of the firm.

Selection is a further problem. Consumers who shop at unusual hours differ from those who do not in ways correlated with income, occupation, and household composition, so cross-sectional differences will confound state with type. This is the same confound the construct is defined to exclude, and it is the reason the definition insists on within-person comparison.

The fairness analysis generates its own empirical exposure. The claim that disclosure of the counterfactual would trigger a fairness reaction is a prediction, not a finding, and the framing sensitivity documented for personalization disclosures cautions that the reaction can be muted by presentation (van Boom et al., 2020). If consumers shown their own within-person price differentials proved indifferent, the third redundancy condition above would be met and the construct would lose its point. Equally, there is a boundary question about which the article takes no position. Some state-conditioned pricing is plainly benign, including a lower price offered to a hesitant consumer, which is the more common commercial practice and which the same inferential machinery supports. A rule addressed to the input catches both. Whether the asymmetric case should be carved out is a policy judgment that

turns on how confidently a regulator can distinguish a discount from a surcharge when only one price is ever displayed.

Finally, the argument would be substantially undermined if within-person price differentials proved negligible in magnitude even where state inference is feasible, or if consumers proved able to detect and defeat them through ordinary comparison behavior. Both are empirical questions, and neither has been answered.

Conclusion

European consumer law contains a serious apparatus for protecting vulnerable consumers, and the informal justice of the market contains a serious apparatus for punishing unfair prices. Both are built on the same premise: that the person at risk can be identified in advance, or that the violation can be perceived after the fact. The premise was a reasonable approximation when commercial practices were addressed to segments. It is no longer a reasonable approximation, because commercial inference has moved to the individual and to the moment, while the protective categories have stayed with the person and the type.

The result is not a single gap but a matched pair of them. The vulnerability premium escapes law because a transient state forms no group, appears on no trait list, distorts everyone's behavior rather than only a minority's, and produces harms too small in each instance to be significant. The one remedy squarely addressed to personalized pricing is a disclosure discharged at the moment of lowest capacity. It escapes fairness norms for the deeper reason that fairness judgment consumes references, comparisons, and counterfactuals, and the practice removes all three. What remains is a transaction in which one party's worst hour is another party's margin, policed by nothing, because the two systems that could police it each wait for a signal the practice is designed not to emit.

For justice research, the premium is more than a policy problem. It is a case in which injustice and the perception of injustice come apart by construction, and it therefore tests the boundary of a field whose instruments are largely perceptual. Studying it requires within-person designs and access to pricing records rather than survey vignettes, and the constructs offered here, the four defining attributes, the decision rule, and the redundancy conditions, are meant to make that study possible and this argument refutable. For regulatory design, the conclusion is narrower. The forthcoming Digital Fairness Act will decide, through a small number of drafting choices about definitions, groups, and disclosure, whether protection continues to track categories of people. Protecting categories remains a coherent ambition. It is no longer a sufficient one, because the market has stopped pricing categories of people and started pricing moments.

Declarations

Author Contribution

S.T. is the sole author. He conceived the argument, conducted the analysis, and wrote the manuscript.

Acknowledgement

The author conceived the research, conducted the analysis, and is responsible for all intellectual content, arguments, and conclusions in this manuscript. Claude (Anthropic) was used iteratively and extensively throughout manuscript preparation for language editing and revision, including substantial rewriting of passages drafted by the author for clarity, structure, and style. The author directed all such revisions, reviewed and approved the resulting text, and takes full responsibility for the accuracy, integrity, and originality of the final manuscript. Every passage the system touched was examined in substance by the author, who holds the subject-matter competence to approve, alter, or reject it and who holds editorial responsibility for the manuscript. The manuscript therefore meets the conditions of the human review and editorial control exemption in Article 50(4) of Regulation (EU) 2024/1689, whose transparency obligations became applicable on 2 August 2026. The declaration is made in full regardless, so that no reader has to infer the extent of the assistance. Machine-readable marking of synthetic output is an obligation of the system provider under Article 50(2) of that Regulation and not of the author; no marking, watermark, or provenance metadata present in any output was removed, altered, or obscured at any stage, and no assisted passage is presented as unassisted. Generative artificial intelligence was not used to generate or analyze data, to produce results, or to supply references. Every source cited was verified by the author against the original publication. No image in this manuscript was produced by a generative artificial intelligence system; the manuscript contains no figures. Claude is not an author of this manuscript and is not listed as one.

References

1. Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–492. <https://doi.org/10.1257/jel.54.2.442>
2. Adams, J. S. (1965). Inequity in social exchange. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 267–299). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60108-2](https://doi.org/10.1016/S0065-2601(08)60108-2)
3. Bar-Gill, O. (2012). *Seduction by contract*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199663361.001.0001>
4. Bar-Gill, O., & Sunstein, C. R. (2025). Algorithmic price discrimination. In *Algorithmic harm* (pp. 15–23). Oxford University Press. <https://doi.org/10.1093/oso/9780197778197.003.0002>
5. Barrafreem, K., & Tinghög, G. (2026). The effect of the Omnibus Pricing Directive on consumer purchase intentions. *Journal of Consumer Policy*, 49(1), 1–19. <https://doi.org/10.1007/s10603-025-09607-z>
6. Ben-Shahar, O., & Schneider, C. E. (2014). *More than you wanted to know: The failure of mandated disclosure*. Princeton University Press.

<https://doi.org/10.23943/princeton/9780691161709.001.0001>

7. Bolton, L. E., Warlop, L., & Alba, J. W. (2003). Consumer perceptions of price (un)fairness. *Journal of Consumer Research*, 29(4), 474–491. <https://doi.org/10.1086/346244>
8. Brennncke, M. (2024). A theory of exploitation for consumer law: Online choice architectures, dark patterns, and autonomy violations. *Journal of Consumer Policy*, 47(1), 127–164. <https://doi.org/10.1007/s10603-023-09554-7>
9. Calo, R. (2014). Digital market manipulation. *George Washington Law Review*, 82(4), 995–1051.
10. Cartwright, P. (2015). Understanding and protecting vulnerable financial consumers. *Journal of Consumer Policy*, 38(2), 119–138. <https://doi.org/10.1007/s10603-014-9278-9>
11. Cole, A. (2016). All of us are vulnerable, but some are more vulnerable than others: The political ambiguity of vulnerability studies, an ambivalent critique. *Critical Horizons*, 17(2), 260–277. <https://doi.org/10.1080/14409917.2016.1153896>
12. Daly, J. P., & Tripp, T. M. (1996). Is outcome fairness used to make procedural fairness judgments when procedural information is inaccessible? *Social Justice Research*, 9(4), 327–349. <https://doi.org/10.1007/BF02196989>
13. Danziger, S., Levav, J., & Avnaim-Pesso, L. (2011). Extraneous factors in judicial decisions. *Proceedings of the National Academy of Sciences*, 108(17), 6889–6892. <https://doi.org/10.1073/pnas.1018033108>
14. Duivenvoorde, B. B. (2015). The consumer benchmarks in the Unfair Commercial Practices Directive. Springer International Publishing. <https://doi.org/10.1007/978-3-319-13924-1>
15. European Commission. (2024). *Commission staff working document: Fitness check of EU consumer law on digital fairness (SWD(2024) 230 final)*. Publications Office of the European Union.
16. Fineman, M. A. (2008). The vulnerable subject: Anchoring equality in the human condition. *Yale Journal of Law and Feminism*, 20(1), 1–23.
17. Folger, R., & Cropanzano, R. (1998). *Organizational justice and human resource management*. Sage. <https://doi.org/10.4135/9781452225777>
18. Frenda, S. J., Berkowitz, S. R., Loftus, E. F., & Fenn, K. M. (2016). Sleep deprivation and false confessions. *Proceedings of the National Academy of Sciences*, 113(8), 2047–2050. <https://doi.org/10.1073/pnas.1521518113>
19. Goodin, R. E. (1985). *Protecting the vulnerable: A reanalysis of our social responsibilities*. University of Chicago Press.
20. Hacker, P. (2023). Manipulation by algorithms: Exploring the triangle of unfair commercial practice, data protection, and privacy law. *European Law Journal*, 29(1–2), 142–175. <https://doi.org/10.1111/eulj.12389>
21. Hanson, J. D., & Kysar, D. A. (1999). Taking behavioralism seriously: Some evidence of market manipulation. *Harvard Law Review*, 112(7), 1420–1572. <https://doi.org/10.2307/1342413>

22. Haws, K. L., & Bearden, W. O. (2006). Dynamic pricing and consumer fairness perceptions. *Journal of Consumer Research*, 33(3), 304–311. <https://doi.org/10.1086/508435>
23. Helberger, N., Lynskey, O., Micklitz, H. W., Rott, P., Sax, M., & Strycharz, J. (2021). *EU consumer protection 2.0: Structural asymmetries in digital consumer markets*. BEUC.
24. Helberger, N., Sax, M., Strycharz, J., & Micklitz, H. W. (2022). Choice architectures in the digital economy: Towards a new understanding of digital vulnerability. *Journal of Consumer Policy*, 45(2), 175–200. <https://doi.org/10.1007/s10603-021-09500-5>
25. Insel, T. R. (2017). Digital phenotyping: Technology for a new science of behavior. *Journal Of The American Medical Association*, 318(13), 1215–1216. <https://doi.org/10.1001/jama.2017.11295>
26. Jasso, G. (1980). A new theory of distributive justice. *American Sociological Review*, 45(1), 3–32. <https://doi.org/10.2307/2095239>
27. Kahneman, D., Knetsch, J. L., & Thaler, R. (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review*, 76(4), 728–741.
28. Killgore, W. D. S. (2010). Effects of sleep deprivation on cognition. *Progress in Brain Research*, 185, 105–129. <https://doi.org/10.1016/B978-0-444-53702-7.00007-5>
29. Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and decision making. *Annual Review of Psychology*, 66, 799–823. <https://doi.org/10.1146/annurev-psych-010213-115043>
30. Lerner, M. J. (2023). Confusing the expression of social norms and justice motivation. *Social Justice Research*, 36(3), 325–335. <https://doi.org/10.1007/s11211-023-00415-9>
31. Leszczynska, M., & Baltag, D. (2024). Can I have it non-personalised? An empirical investigation of consumer willingness to share data for personalized services and ads. *Journal of Consumer Policy*, 47(3), 345–372. <https://doi.org/10.1007/s10603-024-09568-9>
32. Li, H., & Nill, A. (2020). Online behavioral targeting: Are knowledgeable consumers willing to sell their privacy? *Journal of Consumer Policy*, 43(4), 723–745. <https://doi.org/10.1007/s10603-020-09469-7>
33. Lind, E. A. (2023). Focusing on the social in social justice research. *Social Justice Research*, 36(3), 337–351. <https://doi.org/10.1007/s11211-023-00418-6>
34. Lind, E. A., & van den Bos, K. (2002). When fairness works: Toward a general theory of uncertainty management. *Research in Organizational Behavior*, 24, 181–223. [https://doi.org/10.1016/S0191-3085\(02\)24006-X](https://doi.org/10.1016/S0191-3085(02)24006-X)
35. Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, 65(3), 272–292. <https://doi.org/10.1006/obhd.1996.0028>
36. Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, 127(2), 267–286. <https://doi.org/10.1037/0033-2909.127.2.267>
37. Lotz, S. (2015). Engineering fairness? Market design as a resource for social justice research. *Social Justice Research*, 28(3), 391–399. <https://doi.org/10.1007/s11211-015-0250-x>
38. Luguri, J., & Strahilevitz, L. J. (2021). Shining a light on dark patterns. *Journal of Legal Analysis*, 13(1), 43–109. <https://doi.org/10.1093/jla/laaa006>

39. Mathur, A., Acar, G., Friedman, M. J., Lucherini, E., Mayer, J., Chetty, M., & Narayanan, A. (2019). Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1–32. <https://doi.org/10.1145/3359183>
40. McStay, A., & Rosner, G. (2021). Emotional artificial intelligence in children's toys and devices: Ethics, governance and practical remedies. *Big Data & Society*, 8(1), 1–16. <https://doi.org/10.1177/2053951721994877>
41. Mik, E. (2016). The erosion of autonomy in online consumer transactions. *Law Innovation and Technology*, 8(1), 1–38. <https://doi.org/10.1080/17579961.2016.1161893>
42. Mills, S., Costa, S., & Sunstein, C. R. (2023). AI, behavioural science, and consumer welfare. *Journal of Consumer Policy*, 46(3), 387–400. <https://doi.org/10.1007/s10603-023-09547-6>
43. Moya, C., & Adriaans, J. (2026). Justice, inequality, and the joy of discovery: The contributions of Guillermina Jasso to social justice research. *Social Justice Research*. Advance online publication. <https://doi.org/10.1007/s11211-026-00481-9>
44. Nicklin, J. M. (2013). Expertise, counterfactual thinking, and fairness perceptions: A test of fairness theory. *Social Justice Research*, 26(1), 42–60. <https://doi.org/10.1007/s11211-012-0173-8>
45. Poort, J., & Zuiderveen Borgesius, F. J. (2019). Does everyone have a price? Understanding people's attitude towards online and offline price discrimination. *Internet Policy Review*, 8(1), 1–20. <https://doi.org/10.14763/2019.1.1383>
46. Riefa, C. (2022). Protecting vulnerable consumers in the digital single market. *European Business Law Review*, 33(4), 607–634. <https://doi.org/10.54648/eulr2022028>
47. Rogers, W., Mackenzie, C., & Dodds, S. (2012). Why bioethics needs a concept of vulnerability. *International Journal of Feminist Approaches to Bioethics*, 5(2), 11–38. <https://doi.org/10.3138/ijfab.5.2.11>
48. Sax, M., Helberger, N., & Bol, N. (2018). Health as a means towards profitable ends: mHealth apps, user autonomy, and unfair commercial practices. *Journal of Consumer Policy*, 41(2), 103–134. <https://doi.org/10.1007/s10603-018-9374-3>
49. Schuitema, G., Steg, L., & van Kruining, M. (2011). When are transport pricing policies fair and acceptable? *Social Justice Research*, 24(1), 66–84. <https://doi.org/10.1007/s11211-011-0124-9>
50. Seele, P., Dierksmeier, C., Hofstetter, R., & Schultz, M. D. (2021). Mapping the ethicality of algorithmic pricing: A review of dynamic and personalized pricing. *Journal of Business Ethics*, 170(4), 697–719. <https://doi.org/10.1007/s10551-019-04371-w>
51. Shah, A. K., Mullainathan, S., & Shafir, E. (2012). Some consequences of having too little. *Science*, 338(6107), 682–685. <https://doi.org/10.1126/science.1222426>
52. Smith, H. J., Pettigrew, T. F., Pippin, G. M., & Bialosiewicz, S. (2012). Relative deprivation: A theoretical and meta-analytic review. *Personality and Social Psychology Review*, 16(3), 203–232. <https://doi.org/10.1177/1088868311430825>
53. Strycharz, J., & Duivenvoorde, B. (2021). The exploitation of vulnerability through personalised marketing communication: Are consumers protected? *Internet Policy Review*, 10(4), 1–27.

<https://doi.org/10.14763/2021.4.1585>

54. Sunstein, C. R. (2016). Fifty shades of manipulation. *Journal of Marketing Behavior*, 1(3–4), 213–244. <https://doi.org/10.1561/107.000000014>
55. Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, autonomy, and manipulation. *Internet Policy Review*, 8(2). <https://doi.org/10.14763/2019.2.1410>
56. Trzaskowski, J. (2011). Behavioural economics, neuroscience, and the Unfair Commercial Practises Directive. *Journal of Consumer Policy*, 34(3), 377–392. <https://doi.org/10.1007/s10603-011-9169-2>
57. Tyler, T. R. (2023). The organizational underpinnings of social justice theory development. *Social Justice Research*, 36(3), 371–384. <https://doi.org/10.1007/s11211-023-00414-w>
58. van Boom, W. H., van der Rest, J. P. I., van den Bos, K., & Dechesne, M. (2020). Consumers beware: Online personalized pricing in action! How the framing of a mandated discriminatory pricing disclosure influences intention to purchase. *Social Justice Research*, 33(3), 331–351. <https://doi.org/10.1007/s11211-020-00348-7>
59. van den Bos, K., Lind, E. A., Vermunt, R., & Wilke, H. A. M. (1997). How do I judge my outcome when I do not know the outcome of others? The psychology of the fair process effect. *Journal of Personality and Social Psychology*, 72(5), 1034–1046. <https://doi.org/10.1037/0022-3514.72.5.1034>
60. Vrousalis, N. (2013). Exploitation, vulnerability, and social domination. *Philosophy & Public Affairs*, 41(2), 131–157. <https://doi.org/10.1111/papa.12013>
61. Weinshall-Margel, K., & Shapard, J. (2011). Overlooked factors in the analysis of parole decisions. *Proceedings of the National Academy of Sciences*, 108(42), E833. <https://doi.org/10.1073/pnas.1110910108>
62. Wertheimer, A. (1996). *Exploitation*. Princeton University Press.
63. Wieth, M. B., & Zacks, R. T. (2011). Time of day effects on problem solving: When the non-optimal is optimal. *Thinking & Reasoning*, 17(4), 387–401. <https://doi.org/10.1080/13546783.2011.625663>
64. Xia, L., Monroe, K. B., & Cox, J. L. (2004). The price is unfair! A conceptual framework of price fairness perceptions. *Journal of Marketing*, 68(4), 1–15. <https://doi.org/10.1509/jmkg.68.4.1.42733>
65. Yeung, K., & Information (2017). *Communication & Society*, 20(1), 118–136. <https://doi.org/10.1080/1369118X.2016.1186713>.
66. Zuiderveen Borgesius, F., & Poort, J. (2017). Online price discrimination and EU data privacy law. *Journal of Consumer Policy*, 40(3), 347–366. <https://doi.org/10.1007/s10603-017-9354-z>
67. Cases.
68. Case, C. (1998). /96, Gut Springenheide GmbH and Tusky v Oberkreisdirektor des Kreises Steinfurt, ECLI:EU:C::369.
69. Joined Cases C-54/17 and C-55/17, Autorità Garante della Concorrenza e del Mercato v Wind Tre SpA and Vodafone Italia SpA, ECLI:EU:C (2018). :710.
70. Case, C. (2019). /17, Prezes Urzędu Ochrony Konkurencji i Konsumentów v Orange Polska S.A., ECLI:EU:C::480.

- 71. Legislation.
- 72. European, & Union.
- 73. Directive (2005). 2005/29/EC concerning unfair business-to-consumer commercial practices in the internal market OJ L 149/22.
- 74. Directive 2011/83/EU on consumer rights (2011) OJ L 304/64.
- 75. Directive, E. U. (2019). 2019/2161 as regards the better enforcement and modernisation of Union consumer protection rules OJ L 328/7.
- 76. Regulation (EU) (2024). 2024/1689 laying down harmonised rules on artificial intelligence. *OJ L*, 2024/1689.