Allocation Under Fixed Representation

Abstract

Many markets allocate through standardized representations: scores, rating cells, factor categories, catalog taxonomies, or embedding coordinates. This paper studies the consequences of taking such a representation as a fixed action constraint. A full payoff-relevant state may contain more information than the representation records, but admissible instruments are required to be measurable with respect to the representation. Under a quadratic match-to-target benchmark, the optimal constrained action is the conditional expectation of the target given the representation, and the gap to the full-state benchmark is the expected within-cell conditional variance. This is a standard projection identity, not a new variance theorem. The contribution is to use that elementary object as a disciplined diagnostic for a distinct source of allocation loss: information can improve while the action algebra remains fixed. Signal refinement in the Blackwell sense can reduce ordinary information loss, but it cannot reduce the residual associated with a fixed representation. That residual is the loss remaining at common knowledge, not the largest realized loss. The paper also clarifies the limited mechanism-design implication: revelation and implementation results do not help when the final outcome space itself is restricted to the representation, but they may help when the outcome space can be enriched. Applications are developed only where the fixed-representation premise is plausible: standardized credit terms, insurance rating cells, and platform taxonomies or embeddings. The paper deliberately does not claim a general theory of asset-pricing anomalies or a regulation-free explanation of all classification restrictions.

1. Introduction

A central distinction in allocation problems is the distinction between **what the decision maker knows** and **what the decision maker is allowed or able to condition the action on**. These two objects often move together, but they need not. A lender may observe more borrower characteristics than enter the standardized score used to set a conforming loan term. An insurer may learn more about risk than is permitted in a filed rating plan. A platform may log a niche taste that is not represented in the taxonomy or embedding through which ranking, search, advertising, and settlement are implemented.

This paper studies the second object. It asks what follows when the market's instruments must be written on a fixed **representation apparatus**–a score, rating cell, factor category, or taxonomy–rather than on the full payoff-relevant state. The formal assumption is simple. There is a measurable map from the full state to an operative representation. All admissible instruments must be measurable with respect to the sigma-algebra generated by that map. Two states that share the same representation must receive the same instrument value, even if the decision maker observes additional payoff-relevant variation inside the cell.

The mathematical implication is elementary. With a quadratic match-to-target loss, the best representation-measurable action is the conditional expectation of the target given the representation. The resulting gap to the full-state benchmark is the expected conditional variance inside representation cells. That object is familiar from projection theory, coarse matching, exogenous pooling, categorical classification, and restricted-instrument models. The paper does not present the conditional-variance formula as novel.

The point is instead a modeling distinction. If a loss is caused by poor information, then better signals reduce it, and complete information eliminates it. If a loss is caused by a fixed action representation, then better signals may improve beliefs but cannot create a new coordinate on which the action can be written. The residual associated with the representation is therefore invariant to Blackwell refinements of the signal, holding the action algebra fixed. At common knowledge, ordinary information loss has disappeared; the representation residual is the part that remains.

The contribution is deliberately modest and precise. The paper provides a compact model of allocation under a fixed **action algebra**, identifies the resulting projection gap, separates that gap from information loss, and states exactly what mechanism design can and cannot do under the restriction. It also narrows the scope of applications. In credit and insurance, the fixed action algebra may be created by standardized score fields, filed rating plans, legal restrictions, or secondary-market eligibility rules. In platform settings, it may be created by shared taxonomies or embeddings used by multiple platform modules and market sides. These are institutional sources of fixity, not pure information frictions. Where an actor can freely and costlessly change the representation, the premise fails and the model should not be applied.

The paper proceeds as follows. Section 2 relates the argument to coarse matching, price discrimination, insurance classification, restricted menus, incomplete contracts, Blackwell information, and standardization. Section 3 presents the model. Section 4 gives the projection result, the signal-invariance decomposition, the mechanism-design implication, and the monotonicity of representation refinement. Section 5 discusses what it means to hold a representation fixed. Section 6 gives scoped applications. Section 7 states limitations. Section 8 concludes.

2. Related literature

The closest economic relatives are models of **coarse matching** and **exogenous pooling**. McAfee (2002) studies matching when agents are grouped into coarse classes. Hoppe, Moldovanu, and Sela (2009) analyze assortative matching when sorting is mediated by signals. In these and related models, welfare losses arise because agents or objects are pooled into classes that do not fully resolve payoff-relevant heterogeneity. The conditional-dispersion object in this paper is continuous with that literature. The difference is not the variance formula; the difference is the interpretation of the partition. Here the partition is treated as an institutionally fixed action representation rather than as an information partition chosen or refined with the signal.

The paper is also related to **price discrimination** and **market segmentation**. Third-degree price discrimination restricts prices to be constant within observable groups. More recent work on the limits of price discrimination studies how segmentation and information shape feasible pricing. In this language, the present model focuses on the limiting case in which the pricing or allocation partition is fixed exogenously, and the question is what signal improvements can accomplish when they do not alter the partition on which instruments are written.

Insurance classification is another close neighbor. The literature on adverse selection and categorical discrimination in insurance distinguishes risk information, rating categories, and participation responses. The present model isolates only one component: the welfare consequence of rating-cell restrictions when premiums must be constant inside cells. It does not claim that adverse selection is absent in insurance markets, nor that all insurance inefficiency is a representation

problem. Rather, adverse selection and classification restrictions are distinct mechanisms that can coexist.

The model also resembles **restricted-menu** and **restricted-instrument** problems in mechanism design and contract theory. If a principal can offer only a limited menu, or if final outcomes must lie in a restricted set, full surplus may be unattainable even with rich messages. The mechanism-design result below is correspondingly limited: if final outcomes must be measurable with respect to the representation, no message game can implement an outcome that separates states inside a representation cell. This is an outcome-space restriction, not a new implementation theorem.

Incomplete-contracts theory is a different neighbor. Grossman and Hart, Hart and Moore, Moore and Repullo, and Maskin and Tirole study limitations created by noncontractibility, indescribability, unverifiability, and the design of mechanisms under such limitations. The present paper is not an incomplete-contracts model. It assumes that payoff-relevant states may be describable, observable, and verifiable, but that admissible outcomes are still restricted to a coarser representation. If the coarsening in an application is simply the set of legally verifiable attributes, then the model collapses into a standard contractibility restriction. The distinction matters only when the action representation is fixed for reasons other than lack of information or proof.

Blackwell's comparison of experiments provides the language for information refinement. A more informative experiment can reduce decision loss when the feasible action can use the additional information. The present model holds the action algebra fixed. Under that restriction, signal refinement can reduce the information component of loss but not the residual associated with the representation itself.

Finally, the paper connects to work on standardization, compatibility, and commensuration. Farrell and Saloner show how coordination and installed bases can sustain standards even when alternatives are available. Espeland and Stevens emphasize that commensuration organizes some information while discarding other dimensions. This paper uses those ideas only to motivate when a representation may plausibly be fixed. It does not prove a general theory of standard formation.

3. Model

3.1 States, representations, and actions

Let $(Z, \mathcal{S}, P)$ be a probability space of payoff-relevant situations. A situation $z \in Z$ may be a single agent receiving a treatment, a borrower-contract pair, an insured-risk pair, a consumer-product pair, or any other allocation-relevant unit. For pairwise applications, one may take $Z = \Omega \times \Omega$ and $\mathcal{S} = \Sigma \otimes \Sigma$.

A **representation apparatus** is a measurable map

$$\pi : Z \to S$$

from the full state space into a standard Borel representation space S. The induced **action algebra** is

$$\mathcal{A} = \sigma(\pi) \subseteq \mathcal{S}.$$

An action or instrument is **representation-admissible** if it is $\mathcal{A}$-measurable. Thus an admissible action has the form

$$a(z) = \tilde{a}(\pi(z))$$

for some measurable $\tilde{a} : S \to \mathbb{R}$. Two situations with the same representation must receive the same action value.

When regular conditional probabilities exist, one can speak of conditional distributions on the fibers $\pi^{-1}(s)$. These fibers need not be atoms of $\mathcal{A}$, especially in continuous spaces. The formal object used below is conditional expectation with respect to $\mathcal{A}$, which is well-defined in L^2 without treating fibers as atoms.

3.2 The payoff target

Let

$$Y \in L^2(Z, \mathcal{S}, P)$$

be the payoff-relevant target. In different applications, Y may represent an ideal loan term, a risk-adjusted premium target, a match value, or another scalar object that the action tries to track.

The baseline payoff is quadratic:

$$u(a, z) = -(a(z) - Y(z))^2.$$

This is a deliberately special benchmark. It makes the projection structure transparent. For general payoff functions, the welfare gap remains definable, but the conditional-variance formula need not hold exactly.

The unrestricted full-state benchmark allows all $\mathcal{S}$-measurable square-integrable actions:

$$\mathcal{F}_0 = L^2(Z, \mathcal{S}, P).$$

The representation-constrained feasible set is

$$\mathcal{F}_{\mathcal{A}} = L^2(Z, \mathcal{A}, P).$$

The full-state optimum is $a^0 = Y$, with normalized payoff (0). The representation-constrained value is

$$W(\mathcal{A}) = \sup_{a \in \mathcal{F}_{\mathcal{A}}} \mathbb{E}[-(a - Y)^2].$$

Define the **representation gap**

$$R(\mathcal{A}) = 0 - W(\mathcal{A}) = \inf_{a \in \mathcal{F}_{\mathcal{A}}} \mathbb{E}[(a - Y)^2].$$

This is a gap to an unrestricted reference benchmark. If $\mathcal{A}$ is truly fixed by the institutional environment, the benchmark may be infeasible. The gap should therefore be read as a reference measure of the value of the omitted distinctions, not as surplus that was literally destroyed by an available market action.

3.3 Assumptions

The model uses four assumptions.

Assumption 1: Fixed representation. The action algebra $\mathcal{A} = \sigma(\pi)$ is fixed for the problem under study. A change from $\mathcal{A}$ to a finer algebra is a change in representation, not an improvement in predictive accuracy at fixed representation.

Assumption 2: Admissibility. Every instrument relevant to the allocation problem is $\mathcal{A}$-measurable.

Assumption 3: Nondegeneracy. The target Y is not $\mathcal{A}$-measurable. Equivalently,

$$P\left(Y \neq \mathbb{E}[Y \mid \mathcal{A}]\right) > 0.$$

Without this condition, the representation omits no payoff-relevant variation and the representation gap is zero.

Assumption 4: Quadratic benchmark. The conditional-variance formula is stated for the quadratic payoff $u(a, z) = -(a - Y)^2$. For other payoffs, the paper does not claim that the same formula survives unchanged.

The first assumption is the substantive economic premise. The remaining assumptions are technical or definitional.

4. Results

4.1 The projection gap

Proposition 1: Projection under a fixed representation. Under Assumptions 1–4, the optimal representation-constrained action is

$$a^*_{\mathcal{A}} = \mathbb{E}[Y \mid \mathcal{A}],$$

and the representation gap is

$$R(\mathcal{A}) = \mathbb{E}\left[\left(Y - \mathbb{E}[Y \mid \mathcal{A}]\right)^2\right] = \mathbb{E}\left[\mathrm{Var}(Y \mid \mathcal{A})\right].$$

Moreover, $R(\mathcal{A}) = 0$ if and only if Y is $\mathcal{A}$-measurable.

Proof. The conditional expectation $\mathbb{E}[Y \mid \mathcal{A}]$ is the L^2-projection of Y onto the closed subspace

$L^2(Z, \mathcal{A}, P)$. Therefore it uniquely minimizes $\mathbb{E}[(a - Y)^2]$ over all $a \in L^2(Z, \mathcal{A}, P)$. The minimized mean-square error is

$$\mathbb{E}\left[(Y - \mathbb{E}[Y \mid \mathcal{A}])^2\right] = \mathbb{E}\left[\mathrm{Var}(Y \mid \mathcal{A})\right].$$

This quantity is zero exactly when $Y = \mathbb{E}[Y \mid \mathcal{A}]$ almost surely, which is equivalent to Y being $\mathcal{A}$-measurable.

When a regular conditional distribution P_s given $\pi = s$ is available, the same expression can be written as

$$R(\mathcal{A}) = \int_S \mathrm{Var}_{P_s}(Y)\, dP_\pi(s).$$

This is often the most intuitive form: the gap is the average dispersion of the target inside representation cells.

4.2 Information loss and representation loss

A signal experiment generates a signal sigma-algebra $\mathcal{G} \subseteq \mathcal{S}$. A more informative signal corresponds to a refinement of $\mathcal{G}$ in the Blackwell sense. But an information refinement and a representation refinement are different operations.

To separate them, suppose first that the decision rule can condition only on an implementable information algebra $\mathcal{H} \subseteq \mathcal{A}$. The loss from using the best $\mathcal{H}$-measurable action is

$$R(\mathcal{H}) = \mathbb{E}[\mathrm{Var}(Y \mid \mathcal{H})].$$

If $\mathcal{H} \subseteq \mathcal{A}$, the law of total conditional variance gives

$$R(\mathcal{H}) = R(\mathcal{A}) + \mathbb{E}\left[\mathrm{Var}\left(\mathbb{E}[Y \mid \mathcal{A}] \mid \mathcal{H}\right)\right].$$

The second term is an information term: it vanishes as the decision maker learns enough to condition on all of $\mathcal{A}$. The first term is the representation residual: it remains even when the decision maker has complete information about the representation cell, and even when the full state is common knowledge, so long as the action itself must remain $\mathcal{A}$-measurable.

This yields the central diagnostic.

Corollary 1: Signal invariance at fixed action algebra. Holding Y, P, and $\mathcal{A}$ fixed, the representation residual

$$R(\mathcal{A}) = \mathbb{E}[\mathrm{Var}(Y \mid \mathcal{A})]$$

is unchanged by Blackwell refinements of a signal experiment that do not enlarge the admissible action algebra. Such refinements may reduce the information term above, but they cannot reduce $R(\mathcal{A})$.

The residual is therefore not largest at common knowledge. Total realized loss is minimized at common knowledge, because information loss has been removed. The representation residual is the

component that remains after information loss has vanished.

4.3 Mechanisms with restricted outcome spaces

The mechanism-design implication is limited but important. It concerns the final outcome space, not the message space.

Consider a mechanism with message space M, strategy profiles σ, and an induced allocation rule $f_\sigma : Z \to \mathbb{R}$. Call the mechanism **representation-admissible** if every induced allocation rule that can be implemented by the mechanism is $\mathcal{A}$-measurable.

Proposition 2: Mechanisms cannot exceed their outcome algebra. If a mechanism is representation-admissible, then every allocation it implements lies in $L^0(Z, \mathcal{A})$. In particular, no such mechanism can implement an allocation that assigns different outcomes to two states that are indistinguishable under $\mathcal{A}$.

Proof. The conclusion is immediate from representation-admissibility. For every equilibrium strategy profile σ, the induced allocation f_σ is $\mathcal{A}$-measurable. Hence there exists a measurable $\tilde{f}_\sigma$ such that

$$f_\sigma(z) = \tilde{f}_\sigma(\pi(z)).$$

If $\pi(z) = \pi(z')$, then $f_\sigma(z) = f_\sigma(z')$.

This proposition should not be read as a new impossibility theorem in implementation theory. Revelation mechanisms may elicit the full state. Maskin–Tirole and Moore–Repullo mechanisms may be powerful when the outcome space is rich enough. The point here is narrower: if the final outcome itself must be written on $\pi(z)$, then truthful revelation of z cannot be translated into different treatment for states inside the same representation cell. If the mechanism designer can enlarge the outcome algebra, the proposition no longer applies.

4.4 Refining the representation

The representation gap is monotone in the action algebra.

Proposition 3: Monotonicity of representation refinement. If $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{S}$, then

$$R(\mathcal{B}) \leq R(\mathcal{A}).$$

Moreover,

$$R(\mathcal{A}) - R(\mathcal{B}) = \mathbb{E}\left[\mathrm{Var}\left(\mathbb{E}[Y \mid \mathcal{B}] \mid \mathcal{A}\right)\right].$$

Proof. Apply the law of total conditional variance:

$$\mathrm{Var}(Y \mid \mathcal{A}) = \mathbb{E}[\mathrm{Var}(Y \mid \mathcal{B}) \mid \mathcal{A}] + \mathrm{Var}(\mathbb{E}[Y \mid \mathcal{B}] \mid \mathcal{A}).$$

Taking expectations gives the result.

If $\mathcal{B} = \mathcal{A} \vee \sigma(\varphi)$, the difference

$$R(\mathcal{A}) - R(\mathcal{A} \vee \sigma(\varphi))$$

is the value, under the quadratic benchmark, of adding the coordinate φ to the representation. This formula is often the cleanest way to interpret the model. Better information about φ is not enough; the action algebra must be enlarged so that actions can condition on φ.

For non-nested algebras, there is generally no welfare ordering. Two representations can retain different distinctions, and either can dominate depending on the target Y. The relevant order is refinement of action algebras, not the number of variables, the complexity of the model, or the predictive accuracy of an estimator inside a fixed representation.

5. What fixes a representation?

The model takes the fixed action algebra as a primitive. This section clarifies what that premise can and cannot mean.

A representation can be fixed by law, by filed product forms, by secondary-market eligibility rules, by common data standards, by interoperability requirements, by reporting conventions, or by a platform taxonomy used across multiple modules and market sides. These sources are heterogeneous. Some are regulatory. Some are contractual. Some are standardization or coordination constraints. Some may involve opportunity costs. The model does not require these sources to be morally or economically equivalent. It requires only that, for the decision problem under study, admissible actions are functions of π.

This distinction matters because the fixed-representation premise is not secured merely by saying that a richer representation would be inconvenient. If a single decision maker can costlessly append a payoff-relevant coordinate and can then condition all relevant actions on it, the model predicts that the representation gap should fall. In such an environment, persistence of the gap must be explained by something else: optimization, incentives, market power, regulation, computation, or ordinary information frictions.

Nor is the premise secured by relabeling verification limits as representation limits. If the problem is that a court cannot verify an attribute, the model is an incomplete-contracts model in different notation. The present framework is distinct only when the attribute may be observable or verifiable but still cannot enter the instrument because the instrument is constrained to a standardized representation.

The safest interpretation is therefore institutional. The representation is a maintained component of the market architecture for the relevant period and set of instruments. The paper analyzes the allocation consequences of that maintained architecture. It does not provide a general theorem showing that all valuable omitted coordinates will remain omitted.

6. Applications

6.1 Credit scoring and standardized underwriting

Let Z describe a borrower-contract situation, and let Y denote the ideal risk-adjusted loan term under the quadratic benchmark. A credit representation π may include a score, reported tradeline fields, debt-to-income categories, delinquency histories, or other standardized underwriting variables. If loan terms must be functions of π, then two borrowers in the same representation cell receive the same term.

Suppose an additional borrower characteristic–such as income timing, volatility, platform-work history, or a cash-flow pattern–is payoff-relevant within a score cell. If that characteristic is observed but does not enter the standardized representation used for the relevant loan product, then it contributes to

$$\mathbb{E}[\mathrm{Var}(Y \mid \mathcal{A})].$$

Improving the lender's forecast of the omitted characteristic does not change the standardized term unless the term is permitted to condition on a refined representation.

This application should be read narrowly. Some credit restrictions are legal restrictions on admissible variables. Others arise from secondary-market eligibility or standardized reporting. In those cases, the model describes the allocation consequence of the restriction; it does not claim that the restriction is regulation-free or cost-free. If a lender can make portfolio loans using a proprietary richer representation, then the relevant action algebra for those loans is finer, and the predicted representation gap is smaller.

6.2 Insurance rating cells

Insurance provides a direct rating-cell interpretation. Let Y be the actuarially appropriate premium or coverage term for a risk, and let π be the filed rating plan: the variables and categories on which premiums may depend. If premiums must be constant inside rating cells, then within-cell risk variation generates a representation gap.

This is not the same as adverse selection. Adverse selection concerns private information and endogenous participation or coverage choices. A rating-cell restriction can exist even if the insurer and insured both know the risk. Conversely, adverse selection can arise even with rich rating variables if insureds have private information or behavioral responses. The present model isolates the classification component:

$$R(\mathcal{A}) = \mathbb{E}[\mathrm{Var}(Y \mid \mathcal{A})].$$

It says that richer risk information reduces this component only when the rating plan can condition on the richer information.

The model should not be read as a universal claim about insurance markets. Many insurance outcomes reflect moral hazard, selection, administrative costs, legal constraints, and distributional objectives. The representation model is relevant only to the part of the problem created by fixed

rating categories.

6.3 Platform taxonomies and embeddings

A platform may use a representation π of consumers, products, sellers, or content: a catalog taxonomy, a demand embedding, an advertising segment, or a ranking feature map. If ranking, search, advertising, recommendation, and settlement are all implemented through that representation, then actions are $\mathcal{A}$-measurable.

Suppose a niche taste affects match value but is not represented in the platform taxonomy or embedding. If two consumer-product situations have the same representation, the platform's admissible ranking rule treats them alike. The conditional-variance gap measures the match value left unresolved inside representation cells.

This example requires care. A unilateral platform with perfect logs and costless redeployment could change its representation. In that case the fixed-representation premise is not defensible. The model applies instead when the representation is committed for the relevant horizon or shared across multiple platform modules, sellers, advertisers, reporting systems, navigation conventions, or settlement processes. A logged taste then improves the platform's information, but the action remains constrained until the taxonomy or embedding used by the relevant instruments is changed.

6.4 Asset pricing as a boundary case

The model should not be used to infer general asset-pricing anomaly persistence from a vendor factor map. The stochastic discount factor is not generally required to be measurable with respect to a chosen factor representation. If unconstrained traders observe a residual payoff and can trade on it, then the residual may be priced or arbitraged in the usual way.

The projection logic can apply to a narrower mandate: for example, a benchmarked portfolio, a factor-indexed product, or a reporting-constrained intermediary whose admissible positions are functions of a specified factor representation. In that restricted setting, a payoff component orthogonal to the factor action algebra is outside that intermediary's implementable rule. But this is a statement about constrained mandates, not a general theory of equilibrium prices. A market-level asset-pricing claim would require the stronger and usually false assumption that all marginal capital is constrained by the same representation.

7. Limitations

The fixed action algebra is the substantive assumption. The paper does not derive it from primitive preferences, technologies, or equilibrium standard-setting. It analyzes the consequences when such an algebra is institutionally given.

The conditional-variance formula is special to the quadratic match-to-target benchmark. For a general payoff $u(a, z)$, the representation gap is

$$\sup_{a \in \mathcal{F}_0} \mathbb{E}[u(a, z)] - \sup_{a \in \mathcal{F}_{\mathcal{A}}} \mathbb{E}[u(a, z)].$$

That gap need not equal $\mathbb{E}[\mathrm{Var}(Y \mid \mathcal{A})]$. Under smooth concave losses, a local approximation may involve curvature-weighted dispersion, but the exact variance expression is not generally preserved.

The full-state benchmark may be infeasible under the maintained institutional premise. For that reason, the paper uses the term **representation gap** rather than claiming that surplus is literally destroyed. The gap measures the value of distinctions omitted by the representation relative to a reference benchmark.

The mechanism-design result is intentionally modest. It says that mechanisms cannot implement outcomes outside their admissible outcome algebra. It does not challenge revelation principles or implementation theorems when outcome spaces are rich enough to condition on the relevant state.

The measure-theoretic language is also limited. Fibers are useful intuition, but in continuous spaces they need not be atoms with positive probability. Conditional expectations and conditional variances are the formal objects.

Finally, no congruence condition is needed for the projection result itself. If one wants to interpret representation cells as equivalence classes compatible with a trade, matching, or bundling operation, then additional assumptions may be required. When the representation is not compatible with the relevant operation, the correct state space or action algebra must be reformulated.

8. Conclusion

A market can have more information than its instruments can use. When admissible actions must be functions of a fixed representation, payoff-relevant variation inside representation cells remains unresolved by the action, even if it is observed. Under a quadratic benchmark, the resulting representation gap is the expected conditional variance of the payoff target given the action algebra.

The result is mathematically elementary. Its value is diagnostic. If a loss falls with signal quality at fixed action algebra, it is an information loss. If it remains after the relevant state is known and falls only when the representation itself changes, it is a representation loss. That distinction is useful in credit, insurance, and platform settings where standardized scores, rating cells, taxonomies, or embeddings constrain the instruments that can be written.

The model is conservative by design. It does not claim that all persistent misallocation is representational, that mechanisms generally fail, or that asset-pricing anomalies are constitutional. It claims only that, under a fixed action representation, information refinement and representation refinement are different economic operations, and only the latter reduces the residual created by omitted payoff-relevant distinctions.

References

Akerlof, George A. 1970. The Market for Lemons: Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics* 84(3): 488–500.

Armstrong, Mark. 2006. Recent Developments in the Economics of Price Discrimination. In *Advances in Economics and Econometrics: Theory and Applications.*

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. The Limits of Price Discrimi-

nation. *American Economic Review* 105(3): 921–957.

Blackwell, David. 1951. Comparison of Experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press.

Blackwell, David. 1953. Equivalent Comparisons of Experiments. *Annals of Mathematical Statistics* 24(2): 265–272.

Cochrane, John H. 2005. *Asset Pricing*. Princeton University Press.

Crocker, Keith J., and Arthur Snow. 1986. The Efficiency Effects of Categorical Discrimination in the Insurance Industry. *Journal of Political Economy* 94(2): 321–344.

Espeland, Wendy Nelson, and Mitchell L. Stevens. 1998. Commensuration as a Social Process. *Annual Review of Sociology* 24: 313–343.

Fama, Eugene F., and Kenneth R. French. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33(1): 3–56.

Farrell, Joseph, and Garth Saloner. 1985. Standardization, Compatibility, and Innovation. *RAND Journal of Economics* 16(1): 70–83.

Farrell, Joseph, and Garth Saloner. 1986. Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation. *American Economic Review* 76(5): 940–955.

Geanakoplos, John, and Heraklis Polemarchakis. 1986. Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete. In Walter Heller, Ross Starr, and David Starrett, eds., *Uncertainty, Information and Communication: Essays in Honor of Kenneth J. Arrow*. Cambridge University Press.

Grossman, Sanford J., and Oliver D. Hart. 1986. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy* 94(4): 691–719.

Hansen, Lars Peter, and Scott F. Richard. 1987. The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. *Econometrica* 55(3): 587–613.

Hart, Oliver, and John Moore. 1999. Foundations of Incomplete Contracts. *Review of Economic Studies* 66(1): 115–138.

Hoppe, Heidrun C., Benny Moldovanu, and Aner Sela. 2009. The Theory of Assortative Matching Based on Costly Signals. *Review of Economic Studies* 76(1): 253–281.

Magill, Michael, and Martine Quinzii. 1996. *Theory of Incomplete Markets*. MIT Press.

Maskin, Eric, and Jean Tirole. 1999. Unforeseen Contingencies and Incomplete Contracts. *Review of Economic Studies* 66(1): 83–114.

McAfee, R. Preston. 2002. Coarse Matching. *Econometrica* 70(5): 2025–2034.

Moore, John, and Rafael Repullo. 1988. Subgame Perfect Implementation. *Econometrica* 56(5): 1191–1220.

Rochet, Jean-Charles, and Philippe Choné. 1998. Ironing, Sweeping, and Multidimensional Screening. *Econometrica* 66(4): 783–826.

Rothschild, Michael, and Joseph Stiglitz. 1976. Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *Quarterly Journal of Economics* 90(4):

629–649.

Stole, Lars A. 2007. Price Discrimination and Competition. In Mark Armstrong and Robert Porter, eds., *Handbook of Industrial Organization*, Vol. 3. Elsevier.

Varian, Hal R. 1989. Price Discrimination. In Richard Schmalensee and Robert Willig, eds., *Handbook of Industrial Organization*, Vol. 1. Elsevier.

Wilson, Charles. 1977. A Model of Insurance Markets with Incomplete Information. *Journal of Economic Theory* 16(2): 167–207.