2. Robust theory and fragile practice: Information in a world of disinformation Part 1: Indirect communication
Joseph E. Stiglitz and Andrew Kosenko

1. INTRODUCTORY REMARKS AND MAIN MESSAGES

The analysis of economies where information is imperfect and asymmetric has given rise to a revolution in economics (Stiglitz, 2002, 2020). Longstanding fundamental results such as the presumption of the efficiency of competitive markets – Adam Smith's invisible hand – and that competitive equilibrium is always characterized by demand equalling supply have been overturned. Other key results, such as those concerning the existence of competitive equilibrium and its characterization too have been overturned. Since the founding of modern economics, analyses were based on models of perfect information, with the hope that so long as information was not too imperfect, the results would be at least approximately correct. Information economics showed that this was not true: even a little bit of imperfection could drastically change the results.

The insights of this literature have touched virtually every subdiscipline in both macro- and microeconomics – from labour economics to finance, from product markets to insurance markets – and provided intellectual foundations for still other subdisciplines, like accounting and corporate governance.[1] With such an expansive literature, a short survey has to necessarily be selective.[2] We focus particularly on adverse selection, where it is known that there are differences among individuals (among investment projects, among products, among firms), but there is imperfect information about who is who – what Stiglitz (1975a) called *screening.* And rather than discussing the myriad of applications, we focus on some of the general principles and insights that have emerged from the vast literature.

In this chapter we look at the economics of information in instances where information is endogenous – parties form beliefs (subjective probabilities) about the unobserved characteristics of other parties as a result of *actions* taken. There is limited *direct communication*, the one exception being that individuals may disclose *verifiable* information about themselves (their products or their projects). The insights gleaned from the early literature in these settings – the inefficiency of markets, the sharp difference in economic outcomes that incomplete or asymmetric information environments generate relative to complete information environments, as well as the basic mechanisms for overcoming difficulties posed by incomplete or asymmetric information and the nature of welfare improving interventions – continue to be useful in applications and in advancing the theory. At the same time, recent results have modified, and in some cases overturned, key earlier results. In particular, recent literature has reinforced earlier analyses showing the fragility of the results to the precise specification of the information environment. For instance, we note the result of Kosenko, Stiglitz, and Yun (2023) that under quite general conditions, in the absence of communication, no equilibrium exists; both the

price (Akerlof, 1970) equilibrium and the quantity (Rothschild and Stiglitz, 1976) equilibrium can be broken. In the companion Chapter 3, we look at situations where information is instead communicated directly (instead of indirectly, through actions).

Here we survey the advances in the fundamental models of endogenous information – signalling, screening, and adverse selection – that have been made in the decades subsequent to their formulation. The earliest literature *assumed* initial information asymmetries and addressed how they were addressed in competitive markets and monopolies. Here we discuss endogenizing the initial asymmetries, and consider a broader range of mechanisms by which they can be dealt with – including "mechanism design", the most important examples of which are perhaps the design of auctions and matching algorithms.

The earlier literature was written too before we moved into the age of mis- and disinformation. While the sequel chapter deals more explicitly with this issue in the broader context of communication, here we consider explicitly market incentives for obfuscation and the role of public intervention through disclosure and fraud laws.

While market failures are rife in economies with costly information (that is, *all* economies), we explain here how many key so-called reforms of recent years – for instance, "completing markets" through the creation of derivatives and structured finance – have increased systemic informational burdens, undermining decentralization and lowering welfare. Reforms that fail to consider how the proposed changes alter the economy's information structure may well be counterproductive.

The insight that a rational, utility-maximizing decision maker may (and generally speaking, will) reveal her information through her choices and actions – a key insight in different contexts of Spence (1973), Rothschild and Stiglitz (1976), Stiglitz (1977) and Mirrlees (1971) – linked (private) information and (public) actions, and made the analyses tractable. Similarly, Robert Aumann pointed out in the context of two-agent interactions[3] that:

> In the long run, you cannot use information without revealing it; you can use information only to the extent that you are willing to reveal it. A player with private information must choose between not making use of that information and then he doesn't have to reveal it or making use of it, and then taking the consequences of the other side finding it out. (Hart, 2005, p. 694)

Thus, the link between information, action, and inference was made explicit, and "information" became part of the province of standard economic analysis. But critically, information could not be analysed just like an ordinary good, using standard tools, as Stigler (1961) had hoped.

2.   WHY MARKETS WITH IMPERFECT, COSTLY AND ASYMMETRIC INFORMATION ARE NOT EFFICIENT: A BRIEF SUMMARY

A central result of the economics of information is that even seemingly competitive markets are not constrained Pareto efficient – with the term "constrained" emphasizing *taking into account the costs of producing and disseminating information*. Importantly, there are public interventions (using the constrained set of available information) that could make everyone better off. There are several interrelated ideas that help us understand why that is so. Perhaps the most fundamental is that information is a public good – what one individual knows does

not detract from what another does (information is non-rivalrous). The implication of this is that information should be distributed as freely as possible – but doing so would obviously not be in the interests of any individual or firm with information they believe to be valuable and not widely known.

Nonetheless, it may be difficult to exclude some individual from the benefit of information provided by another.[4] If informed individuals buy more of a stock because they know it will do well, the increase in the price of the stock conveys information from the informed to the uninformed. This makes it impossible for those expending resources to obtain the full social return from those expenditures. There will be underinvestment in research finding out about what are good investments, reflecting the general principle: Markets on their own are *never* efficient in the provision of a public good. More generally, welfare losses are associated with the incomplete appropriation of the benefits of "information externalities" (Stiglitz, 1975b,[5] 1975c, and Leitzinger and Stiglitz, 1984 provide a specific application in the context of oil exploration) and the pecuniary externalities that are pervasive when there is imperfect and asymmetric information (Greenwald and Stiglitz, 1986).

Whenever information is costly to obtain and transmit, there is imperfect and asymmetric information, and whenever that is true, there may not be a full set of markets – certain markets may be "shut down", in particular key risk markets.[6] But while Arrow (1964) showed that the existence of a full set of markets was a sufficient condition for Pareto optimality, subsequently Stiglitz and others (Stiglitz, 1982a; Greenwald and Stiglitz, 1986; Geanakoplos and Polemarchakis, 1986) showed that generically, whenever there were not a full set of markets, the market economy is not constrained Pareto efficient, i.e. having a full set of markets was essentially *necessary* for Pareto efficiency.[7] But "completing the market" by adding markets may be welfare-decreasing if it is not complete – just another application of the theory of the second best.[8]

More generally, whenever there are not a full set of markets, whenever there are incentive compatibility and self-selection constraints (which help screen individuals and ensure that they provide effort),[9] and whenever there is costly search, there are pecuniary externalities that matter (Greenwald and Stiglitz, 1986, 1988; Arnott et al., 1994). Individuals assume, for instance, that health insurance premia are unaffected by how much they smoke, but when *all* individuals smoke, premia do increase. These results overturn a central pillar of standard economics, the first fundamental welfare theorem, the formalization of Adam Smith's conjecture that the pursuit of self-interest would lead, as if by an invisible hand, to the well-being of society.

The Greenwald-Stiglitz theorem establishing the *constrained* Pareto inefficiency of the market put a new light on one of the central ideas in economics, the use of the price system for decentralization. It was obvious that once it was recognized that observable actions affected information, individuals might change their actions to convey to others that they were, for instance, more able or less risky than they really were, and that others were. But *given the imperfections of information*, and that there would be costs, one way or another, in differentiating individuals (in the case at hand, through self-selection constraints, whether in a screening or signalling model), it was not obvious that the competitive market would not be constrained Pareto efficient nonetheless. Indeed, in some simple early examples involving only one commodity (such as the Rothschild-Stiglitz world), the economy was constrained Pareto efficient. But once one went beyond that special case, it was not in general efficient. Interactions *across* markets mattered, so unfettered decentralization is not efficient and designing interventions

that would enable efficient decentralization is very difficult; and when those cross-market externalities are particularly important, as in agrarian markets with sharecropping, decentralization simply breaks down.[10] (See section 4 for a further discussion of decentralization.)

While the focus of this survey is general theory/microeconomics, it should be obvious that these market failures have important implications for macroeconomic performance, including macroeconomic externalities (Korinek, 2012; Jeanne and Korinek, 2010; Davila and Korinek, 2017), financial frictions (Bernanke and Gertler, 1989; Greenwald and Stiglitz, 1993; Stiglitz and Greenwald, 2003; Stiglitz and Weiss, 1992), and labour market rigidities (Delli Gati et al., 2012).[11]

As we just observed, the dimensionality of the price system *within existing markets* is lower than the dimensionality of relevant information, so prices alone cannot in general convey all the relevant information, and existing prices may do "double duty", conveying information not just about scarcity. Much of the formal discussion below is concerned about how, in such situations, *additional information* (say, about quantities) conveys further information, overturning one of the central results of standard economics that prices convey *all* the relevant information. And the fact that that is so not only changes the economy from what it would look like in a world where, say, quantities did not convey information, but even results in the economy not being constrained Pareto efficient. These results also overturn another pillar of standard economics, the notion that markets are informationally efficient, conveying all the relevant information from the informed to the uninformed.[12]

There is still one more reason that markets with imperfect and asymmetric information are not constrained Pareto efficient: a central insight of the earlier information literature was that competitive market equilibria may be characterized by markets not clearing.[13] Whenever that is the case, shadow prices (say, of capital in the presence of credit rationing) will not equal market prices, and not surprisingly, market allocations will again not be Pareto efficient.[14]

Finally, markets with imperfect information are likely to be imperfectly competitive, for a whole variety of reasons. The fact that it is costly to search gives firms market power over their customers and workers.[15] The fact that current employers have more information about their employees than others creates an impediment to labour mobility, giving even more market power to employers (Greenwald, 1979, 1986). Because of space limitations, we do not discuss this important source of market failure further in this chapter.

These market failures are critical to understanding behaviour in many of the key markets in the economy (insurance markets, financial markets, labour markets), which also have large macroeconomic consequences, and obviously important policy implications.

3.   IMPERFECT AND ASYMMETRIC INFORMATION AND IMPLICATIONS FOR MARKETS

3.1   Voluntary, Truthful Disclosure and the Walras Law of Screening

There are some settings where information may be verified (perhaps by an independent third party, such as a notary, a mechanic in the car example, an evaluation by an outside healthcare professional in the insurance example, or a standardized test in the education example). What happens when information is imperfect and asymmetric, but *verifiable*, in the sense that there

is something the imperfectly informed party can do to confirm (or disconfirm) the information, or find out the quality or type?

The central mechanism driving outcomes in this setting was first observed by Stiglitz in 1975, when he wrote:

> […] assume the most able is able to provide information certifying to his abilities. The market would then, in equilibrium, pay the remaining workers their (now lower) mean marginal productivity. It would clearly pay, then, for the most able person of this group to have his ability certified. And the analysis proceeds, until information about the capabilities of all individuals except for the least capable is provided: but if we have sorted out all except for the least capable, we have also sorted out the least capable. This may be called the Walras Law of screening information. (Stiglitz, 1975a, p. 287)[16]

This key result, also known as the "unravelling" result, says that under certain conditions, with verifiable information, private information gets fully revealed in equilibrium (Stiglitz, 1975a). It always holds if it were costless to get a credential verifying one's ability. This "Walras law of screening" was subsequently studied by Milgrom (1981), who confirmed that this was the unique equilibrium when information and verification is costless. Thus, the conclusion of this literature (known as the *verifiable disclosure* or *hard information* literature) is that if information is costlessly verifiable, in equilibrium it always gets revealed.

Market outcomes in which different groups are not differentiated (because of a lack of information) were referred to by Rothschild and Stiglitz (1976), and Stiglitz (1975a), as "pooling" equilibria. Stiglitz (1975a) showed that if verification is costly, there may not be full information revelation; there may also be a pooling equilibrium. Only the more able individuals benefit from the validation, and so they will expend the resources to do so if the costs of screening are small enough. In the pooling equilibrium, the wages of the more able reflect the *average* ability. If the costs of verification are high enough, the difference between the wage they get if they show they are more able – their true productivity – and that average is smaller than the cost of verification, so a pooling equilibrium can be sustained. Thus, in this economy with costly verifiable credentials, there may be multiple equilibria – both a pooling equilibrium (where no one's ability is revealed) and a separating equilibrium (where everyone's ability is revealed). Interestingly, in this case with multiple equilibria, everyone in the pooling equilibrium is better off than in the "separating" equilibrium; still, the Pareto inefficient equilibrium can be sustained.[17]

Much of the subsequent literature has explored situations where these various forms of equilibrium may arise: pooling, separating, and hybrid (partially pooling/partially separating); whether a competitive equilibrium exists at all; and analysed the welfare economics of each pattern of equilibrium. While screening with verifiable information generated multiple equilibria, including a pooling equilibrium, in models with screening with self-selection (Rothschild and Stiglitz, 1976), there was no competitive equilibrium if the differences among individuals were small, and there *never* existed a pooling equilibrium. (We elaborate on this below in subsection 3.10. See also the discussion of repeated interactions, subsection 3.9 below.)

The welfare analysis entails ascertaining the differences between private and social costs and benefits of differentiating. For instance, in the context of screening/signalling, social returns are related to the better resource allocations that are generated by better information – which may be nil;[18] while the private returns include the increased rents associated with being identified as better (more able), in the context of labour markets, having higher returns

in the context of investment markets, higher probability of repaying, in credit markets, or, in insurance markets, a lower probability of the insured-against event occurring. The simple appropriation of ability rents increases inequality without increasing productivity.

But these results with verifiable disclosure of information contrast markedly with those with self-selection. There, the equilibrium, if it exists, is Pareto efficient;[19] but when there exists no equilibrium, Pareto efficient outcomes can be sustained with cross subsidies from the policies purchased by low-risk individuals to those purchased by the high risk (Stiglitz, 2009).

3.2 Different Structures of Information Asymmetries and the Existence of Competitive Equilibrium

In the models discussed so far, the more able individuals know they are more able. There are, however, some contexts in which that is not the natural assumption. For instance, life insurance companies may have better information about the correlates of longevity than ordinary individuals. In that case, again it is the more informed party that pays for the screening: it is advantageous for the life insurance company to screen individuals to determine which individuals are likely to have higher longevity – so long as that information does not become public. But if it does – if others can see that the insurance firm has offered the individual insurance at a low price (reflecting low risk), then other companies will do so also, and the firm doing the screening would be unable to appropriate the returns to its investment in screening. Thus, a screening equilibrium can exist only if markets are (as in Grossman and Stiglitz, 1980) at least partially informationally inefficient, i.e., only if the actions of the screener are only partially observable. Emran and Stiglitz (2009) show that this insight, the inability to appropriate returns from screening, may play an important role in inhibiting lending to young entrepreneurs.

Moreover, there are problems in sustaining a competitive equilibrium when firms do the screening, even if the information they obtain doesn't leak out. Indeed, if even two firms uncover the information about the skills of a particular individual, for example, Bertrand competition amongst them will result in all of the gains going to the individual and the firms that have expended resources to screen will lose money. With no one screening, it would pay a single firm to screen. But if two firms screen, everyone loses money. The only equilibrium is a mixed strategy equilibrium, where firms randomly choose who to screen, with some individuals then being screened only once while others have their abilities identified by multiple firms (Stiglitz, 1975b). That raises several questions, to which we now turn.

3.3 Endogeneity of Information Asymmetries: Learning About Oneself

First, if initially there are no asymmetries of information, but individuals know that there are differences in individual abilities,[20] would it pay them to first identify their abilities for themselves, and then, if they are able, spend still more money verifying this for third parties? From a social perspective, such expenditures are problematic because if there were no productivity benefits from such information, *ex ante*, behind the veil of ignorance, all individuals are better off in the pooling equilibrium, as we have already noted. Yet Stiglitz (1984) shows that there exists an equilibrium in which nonetheless everyone pays to find out their abilities and then the more able provide verification. This is socially unproductive and simply increases inequality. On the other hand, if different individuals have different *comparative* advantages,

then the information will be socially productive. Still, there is no presumption that the market equilibrium is efficient even when there are *some* social gains from better resource allocation, and in particular, there is no presumption that those gains outweigh the social welfare loss of increased inequality.

Another strand of the endogenous information acquisition literature explores when one wouldn't acquire a costly signal (see Sims, 2003; Matějka and McKay, 2015; Caplin and Dean, 2015; Caplin et al., 2022; Maćkowiak et al., 2019).

3.4   Excessive Investment in Information

The excessive investment which creates the asymmetries of information compounds the inefficiency noted earlier that arises in providing verifiable information about ability differences. This is not the only situation in which equilibrium may be characterized by excessive information, where there are social costs to obtaining information that are incommensurate with the benefits – indeed, the social benefits to an abundance of information may be negative (Hirshleifer, 1971). Insurance markets are particularly fragile: disclosure of genetic information may lead to the unravelling of certain insurance markets[21] (see Rothschild and Stiglitz, 1997). In some cases, banning the use of certain information may be Pareto efficient (this may be the case for anti-discrimination laws which move the economy *costlessly* from a discriminatory equilibrium to a Pareto superior non-discriminatory equilibrium (Stiglitz, 1973, 1974b)); but in other cases, such interventions lead to the use of correlated variables, resulting in an even less desirable equilibrium (Rothschild and Stiglitz, 1982).

There is another situation where too much information is "extracted": that where firms have some market power. Such firms have an incentive to acquire information that may enhance the ability of the firm to price discriminate, to extract more of the potential consumer surplus from its customers. Extracting the consumer surplus is not simply distributive, with money often going from poorer consumers to the richer owner of firms; it is a costly "adverse distribution". This rent extraction can be a major source of the firm's profits as well as the distortions associated with market power (Stiglitz, 1977).

3.5   Creating Information Asymmetries

In some cases, economic agents may deliberately *create* information asymmetries.[22] For instance, Edlin and Stiglitz study a setting where firm managers choose to invest in riskier (i.e., noisier) projects because it increases perceived uncertainty about the firm's prospects, and therefore discourages either other firms taking over the firm or competing managers from displacing the current manager (Edlin and Stiglitz, 1995). It is not just the uncertainty that matters; it is that outsiders know less (have less precise information about the firms' assets) than the manager. Thus, the level of informativeness of the decision maker, and the level of informativeness of others, is determined endogenously, and in ways which may be suboptimal.

3.6   Disclosure Requirements

While the analysis of sections 3.3 and 3.4 suggested that there may be situations where in equilibrium there is, in some sense, too much information, public policy has been concerned with cases where, without government intervention, there is too little disclosure. In credit markets,

lenders are required to disclose their true effective interest rate. The US Consumer Financial Protection Bureau (CFPB) likewise requires lenders to state their terms transparently. Here, the concern is that lenders present information in ways in which it is not fully understood by a large fraction of borrowers. There is an attempt to deceive. The US Securities and Exchange Commission (SEC) regulations typically require firms to disclose truthfully all materially relevant information about the securities they are issuing. These regulations go beyond requiring "the truth, nothing but the true, but not necessarily the whole truth". Implicitly, they take the view that knowing that there is a serious downside risk to an investment and not disclosing it is effectively a lie. Their perspective rejects *caveat emptor*, which effectively puts all the burden of information on the buyer. If the seller knows something that he reasonably should know would be relevant for the buyer, he must disclose.

There are several justifications for such disclosure requirements and several contexts in which there will be insufficient voluntary disclosure. A strong assumption in the above analyses yielding full disclosure is that of rational expectations – that the uninformed party understands the distribution of the qualities of the products being purchased, given the (limited) information being provided; she simply doesn't know which item is of which quality. Moreover, she rationally makes inferences from what is observed. Behavioural economics has taught us that that is not the case. Cognitive limitations may lead individuals to make the wrong inferences, and firms have learned how to better "deceive" individuals. There are extensive literatures documenting persistent, systematic biases in consumer behaviour (including incorrect estimates of own future behaviour); the fact that these biases are exploited by firms can lead to subpar choices with regard to health, borrowing, saving, and investing behaviour.[23] Providing *standardized* disclosure statements enables better and less costly assessments of the relative merits of different products or investment opportunities.

Moreover, as we comment further below, providing only partial information imposes costs on other market participants. Some may have such high costs of screening that they simply accept the randomness of the product. But others with low costs of screening differentiate, but these are costs they would not have to bear in the presence of required disclosures.[24]

An analogy may be useful. In the economics of liability, Calabresi (1970) argued that the burden should be placed on the party that could avoid an accident at lowest cost. Here, *not* disclosing imposes significant costs on others, either in terms of making a suboptimal product selection or imposing costs of obtaining information – costs that could easily be avoided or reduced if there were disclosure.[25]

We noted above the importance of standardization. Standardization in the way that information is presented is important in a world with a superabundance of information: firms who do not want to make disclosures would otherwise comply with disclosure rules by burying the relevant information in a large volume of irrelevant information designed to overwhelm the attention and information processing capacities of the reader.[26] It is important that information not only be disclosed, but be disclosed in a way that can be analysed and interpreted at relatively low cost. This is a form of signal jamming which is discussed more extensively in the next chapter.

3.7   Fraud Laws

If there were no costs to verifying, as we have noted, then the equilibrium that emerges is one in which everyone truthfully discloses. But if there is a cost to verification, and if everyone

were telling the truth, it wouldn't pay anyone to spend the upfront cost of verification. Hence, an equilibrium with full information couldn't be sustained. There are several ways around this conundrum.

One is that the individual provides a guarantee that his statement is truthful, giving the money back (plus the cost of enforcement of the guarantee plus an arbitrarily small additional cost).[27] Then, no one would have an incentive to lie. (A guarantee, of course, is not just a signal; it is also an insurance policy. Even if there were no asymmetries of information, the seller might provide insurance because it is in a better position to absorb the risk of the non-performance of the product. The fact that it thereby has an incentive to make the risk of non-performance lower is also important.)

An alternative is criminalizing fraud – punishing untruthful statements with a penalty at least equal to the amount received.[28] Not disclosing materially relevant information, particularly information which is easily at hand, may, as we have noted, be viewed as akin to fraud.

There is a remarkable paucity of literature on fraud, in spite of its importance (but see Greenwald and Stiglitz, 1992). Indeed, fraud played a critical role in the financial collapse of 2008 (Stiglitz, 2010).

Since with a guarantee, there still has to be an arbiter of "truth" (is the product what the seller claims it to be?) and the guarantee may have to be enforced through a court of law, the basis of "truthfulness" in both cases is legal enforcement. This is where the assumptions of verifiability and enforceability become central. In the general case, information is only imperfectly or imprecisely verifiable, particularly to third parties. On the other hand, in the context of repeated interactions (repeated games), reputation may serve as an enforcement mechanism and third-party verifiability may not be required. See the brief discussion in section 3.9 below.

The issues under discussion here become particularly relevant in a world of mis- and disinformation – a world into which we seemed to have descended. The early literature, for the most part, while recognizing information asymmetries, assumed that when information was disclosed, it was truthful. The discussion in the next chapter explores these issues further.

3.8   Costly Signals and Self-Selection

In many cases, rather than the costly acquisition of a verified credential, market participants look toward a "surrogate" for the credential, a costly signal that is correlated with the relevant, unobservable ability or quality, and a major strand of the literature has focused on this (Rothschild and Stiglitz, 1976, hereafter referred to as RS, and Spence, 1973).

Sometimes the surrogate is simply a more easily observable characteristic: if race is correlated with education, and education is correlated with ability, but race is more easily observable than education, race may be used to screen individuals, resulting in a discriminatory equilibrium (Phelps, 1972; Arrow, 1973; Rothschild and Stiglitz, 1982; Stiglitz, 1973, 1974b).

Most of the literature has focused though on cases where the surrogate is a decision variable, and knowing that a change in action may affect the inferences about, say, ability has fundamental effects on the nature of market equilibrium. The better-informed party may try to signal his qualities. Or the uninformed party may structure a set of choices where the choices *reveal* who is better: self-selection constraints work to enable the firm to "screen". When there exists a separating equilibrium (where there is full disclosure, the "unravelling" equilibrium discussed in the previous section), a self-selection (RS) equilibrium and a signalling equilibrium are typically the same.

The array of actions that convey information – and therefore may well be affected – is enormous: the quantity of insurance as in RS, the level of education, as in Spence, the number of hours one is willing to work a week, as in Akerlof (1976), the willingness to engage in search, as in Salop (1977). A large literature has developed exploring the mechanisms that might be used in different contexts. A key implication was noted in the beginning of the chapter: virtually *all* behaviour is affected in one way or another.

In the signalling model (Spence, 1973), the informed party moves first, taking a costly action (such as a level of education), which is a signal of the attribute or quality in question (in the context of education, the individual's ability). Although expenditures on the signal are assumed not to increase the individual's productivity – thus signalling is purely dissipative – the action verifiably differentiates her from others. In the screening models (RS), the uninformed party moves first, structuring a set of choices (actions) that differentiates among individuals.[29]

The two strands of literature (testing/verification, signalling/self-selection) are actually more closely related than has been recognized. Assume that anyone can pass the test for "high ability" *if they put in enough effort.* Thus, passing the test is a noisy signal of ability, confounded by another unobserved variable, effort. But if the test is hard enough, then it would not pay the low-ability person to exert the effort required to pass the test. Passing the test of sufficient severity thus becomes a costly signal. The cost of the signal – the cost of information – is thus both the cost of the test and the effort to pass the test.

Of course, sometimes the "action" has value in its own right – getting more education increases productivity. Then the social cost of the screening/signalling information is the expenditure on education *beyond* that which would be desirable in a world of perfect information.[30,31]

Much of the theoretical work in signalling models has focused on how the different parties reason upon observing signals that should not be used, that would not be observed in equilibrium. The literature on "refinements" that operates by restricting agents' beliefs (Banks and Sobel, 1987; Cho and Kreps, 1987), and the rest of the literature that operationalizes the "stability" notion of Kohlberg and Mertens (1986) has provided a veritable bestiary of "equilibrium refinements" that aim to obtain "reasonable" outcomes in signalling games. They proceed axiomatically, stating desiderata for what good (what they refer to as a "stable") equilibrium (among all the possible equilibria) looks like. For instance, a "good" outcome in the Spence signalling game is the least-cost separating equilibrium[32] (the "Riley outcome", Riley, 1979)[33] where the high types obtain as little education as necessary (and thus incur as low a cost as possible) to distinguish themselves, and there is information revelation by virtue of the different ability types choosing different levels of education, and thus signalling their types to the employers. In other words, there is information revelation at the lowest possible cost to the participants.[34]

3.9   Multiple-Period and Repeated Interactions, Reputation, and Robustness

Even when characteristics cannot be verified by a third party – and therefore guarantees and the truthfulness of statements cannot be enforced – unravelling (a separating equilibrium) may occur through reputation mechanisms. In the *reputation* literature, repeated interactions allow some *observability* (without third-party verifiability), but for such mechanisms to work, outcomes or characteristics must be observed with sufficient precision (or signals of charac-

teristics or actions, which are sufficiently correlated with outcomes or characteristics, even if imperfectly).[35] In the absence of third-party verifiability, the main enforcement is through a cut off in relationships,[36] which itself can be costly to both parties.

Equilibria in reputation models with, say, unobservable effort have one distinct characteristic which distinguishes them from standard competitive equilibria: there have to be "rents", with prices, for instance, in excess of costs. What induces good behaviour is the threat of losing those rents (Shapiro, 1982; Shapiro and Stiglitz, 1984) as a result of a cut off in the relationship.

Screening and signalling in multi-period contexts also change incentives in a fundamental way: if an individual reveals who he is in one period, that information can be used in all subsequent periods. It increases the incentive of the more able to have himself distinguished from the less able, but also increases the incentive of the less able not to have himself distinguished from the more able. The latter effect may dominate: in a model with a finite number of periods, there may exist a pooling equilibrium for an initial set of periods (Gale and Stiglitz, 1989; Courty and Hao, 2000; Stiglitz, 2009).[37]

The literature on repeated games has also investigated how robust the predictions are to a misspecification of the problem (for instance, if the beliefs about the preferences, information, or behaviour of one or more of the agents are incorrect). Hörner, Ely and others have worked on so-called "belief free" equilibria (Ely et al., 2005; Kandori and Obara, 2006; Hörner and Lovo, 2009; Hörner et al., 2018), equilibria that are completely independent of any information held by others, and rely only on individual optimization, and are accordingly extremely robust to informational assumptions.[38] A profile of action sequences (one for each player) in a repeated game is *belief-free* if after any history of actions (some parts of which may be unobserved), a player's plan of action from that point onward is a best response to the plans of action of the other players that are in the profile of action sequences we started with, for *any beliefs* the player may hold about others' plans of action (or, equivalently, regardless of any private information a player may have). In a belief-free equilibrium of a repeated game, at every point every player's strategy of play from that point onward is optimal, given her information, *independently* of the information held by the other players. They are, by definition, robust to (in fact, completely independent of) how individuals form their beliefs. This is an extremely strong assumption. Technically speaking, it is a subgame perfect equilibrium for every game of complete information that is consistent with the player's own information. Thus, players need not use Bayesian reasoning about others and adhere to the Harsanyi doctrine (discussed below). Of course, the assumption of individual optimization in complex environments (and even of maximizing *expected* utility) itself is an extremely strong requirement (see the discussion in the next chapter on behavioural economics).

A small diversion into the theory of repeated games may be warranted. There exists a strong result in the theory of repeated games that states the following: take the stage game (the game repetitions of which will become the repeated game). Find equilibria of this stage game, the corresponding payoffs, and now take the so-called convex combination of those payoffs (that is, if in one equilibrium a player earns a payoff of 1, and in another equilibrium she earns a payoff of 2, a convex combination is any number between 1 and 2), and suppose that for each player, they are above a lower bound (a payoff that a player can *guarantee* herself, for any set of actions of her opponents). The *folk theorem* states that *any* such payoff (a convex combination of equilibrium payoffs that is above a lower bound) can be a payoff in the repeated

game, provided players are patient enough. Thus, the repeated game has incredibly large sets of payoffs, and therefore, even larger sets of equilibria.

A typical reasoning for why the folk theorem holds goes as follows – suppose a player deviates while playing a game, taking an action that gives her higher instantaneous utility, but lowering the payoffs of others. It is then in the best interest of the other players to play something that "hurts" the deviating player enough (i.e., for long enough), that the deviator will not cheat in the first place. In fact, there is a family of such results, which vary the equilibrium concept (Nash, subgame-perfect, sequential, etc.), and vary what happens after a player deviates – can there be forgiveness, what happens if the punisher(s) deviates, how to evaluate infinite streams of payoffs, and so on.

One important consequence of belief-free equilibria in repeated games is that typically, the folk theorem fails to hold if belief-free equilibrium is used as a solution concept: it's not true that nearly any feasible payoff that is above a lower bound can be the outcome in a belief-free equilibrium. In fact, a belief-free equilibrium may not exist at all. The definition of belief-free equilibrium seems to be too strong to sustain *any* feasible equilibrium payoff, and while it may sustain some (even a "large" set of payoffs), this is still typically smaller than the set of payoffs sustained by the folk theorem.

3.10   Equilibrium – or Non-Existence of Equilibrium – Under Different Information Structures

A key determinant of the nature of the equilibrium is what is observable and what can be communicated to whom and by what means. For instance, the adverse selection problem long investigated in the insurance market assumed that quantities of insurance were not observable; market participants only know that the average quality of what is being traded on the market (the average risk of those willing to buy insurance) is affected by the price.[39] If, as in Rothschild and Stiglitz (1976), quantities of insurance purchased are observable, then that information conveys information about the individual's type: now it is not just prices but quantities that are informative. Thus, the nature of the equilibrium is highly sensitive to whether an individual's purchases of insurances is observable.

But both the adverse selection and the RS equilibria are fragile. As Kosenko et al. (2023) (or KSY for brevity) point out, a single insurance company *can* obtain some information about "quantities" – even if total purchases are not observable – simply by selling a large insurance policy. KSY establish that such limited information is enough to break the standard price equilibrium. Similarly, RS had shown that with full observability of quantities purchased there exists an equilibrium in a model of two types if the two types differ by enough. But if we now modify that world to include the possibility of "secret" insurance (i.e., the possibility that *some* insurance contracts are not observable), then there *never exists an equilibrium* – neither a price equilibrium, nor a quantity (RS) equilibrium, nor a price-cum-quantity equilibrium of the kind that we just described as breaking the price equilibrium.

The disturbing implication of KSY is that in this more realistic world where there is some but not full observability of insurance purchases, an equilibrium never exists. We'll see in the next chapter that if there can be direct communication between consumers and insurance firms and insurance firms with each other, this conundrum is resolved: there always exists an equilibrium, and it entails a pooling contract.

In practice, exclusivity cannot be enforced. In the context of insurance, there are a large number of informal implicit risk sharing mechanisms; risk is, for instance, shared within families. A natural question is: are such informal risk sharing mechanisms welfare improving? At one time, there was the hope (and it was just that, a hope) that social institutions might "fill the gap" arising from the absence of markets or other market failures, reducing the need for government intervention. For instance, concerns about moral hazard limited the amount of insurance that the market would provide (with complete insurance, individuals would take no care). Arnott and Stiglitz (1991) showed, however, that these non-market institutions, *which naturally arose to fill the gap*, can lower welfare, crowding out more efficient market insurance.[40]

3.11   Information Extraction in Monopolies and with Imperfect Competition

The discussion so far has focused on how in a *competitive* world with asymmetric information, those information asymmetries are (partially) overcome, sometimes at great cost, and typically affecting in fundamental ways the nature of the equilibrium. In many ways, the problem of overcoming information asymmetries is simpler for a monopolist: he alone constructs the choice set; he doesn't have to worry about another firm making an offering. With perfect information enabling perfect price discrimination, a monopolist could and would extract all consumer surplus – and there would be no monopoly distortion. Stiglitz (1977) shows how a monopolist can maximize rent extraction (i.e., capturing as much of the consumer surplus as possible, *given the limitations in his information*) in a situation where it knows different consumers differ, but can't tell which are the ones with high consumer surplus. Most importantly, as we noted earlier, contrary to standard analyses, this attempt to differentiate among consumers is the real source of monopoly distortion.

Salop (1977) provides a telling example, where a monopolist charges different prices in different stores – forcing unnecessary search – simply to enable him to charge a higher price to consumers with high search costs, who in his model have higher consumer surplus, increasing thereby the rents he can extract.

In the next chapter, we explain how these problems may be exacerbated with the digital platforms and artificial intelligence, in ways which fundamentally undermine the efficiency of the market economy.

Markets with monopolistic competition and oligopolies lie somewhere between the two polar cases of pure monopoly and competition: firms may have to be sensitive to the offerings of others (and therefore, of the information it can extract from its own offerings alone), but have greater discretion than they do in highly competitive markets, where, for instance, an attempt to "cream skim" the best customers is met with a strong competitive response. Not surprisingly, the equilibrium in general entails distortionary screening mechanisms, including price discrimination and price dispersions. When, for instance, individuals differ in their search costs, market equilibrium will be characterized by price dispersions, inducing unnecessary search, in an attempt to extract more money out of the high search cost to individuals, not unlike that which occurs under monopoly. Indeed, the only market equilibrium may be one with price dispersion, *even when all firms are initially the same* (Salop, 1977; Salop and Stiglitz, 1977, 1982).

3.12   Information Extraction by Governments

In the previous subsection, we saw how monopolists attempt to differentiate among customers in order to extract more of their consumer surplus. Governments too want to differentiate: equalitarian governments would like to tax more those who are more able, to redistribute to those who are less able. Regulators want to ascertain whether a utility should be allowed to charge a higher price, because its costs are genuinely higher. There is a close formal similarity between the solution to the problem of maximizing social welfare by a government or regulator, in the presence of asymmetries of information, and the problem of a monopolist maximizing its profits, in the presence of asymmetries of information.[41]

3.13   Summary of Standard Literature

Five fundamental insights that have emerged in the voluminous literature of the past nearly half century are (a) there may exist no equilibrium;[42] (b) while in certain classes of models, the only equilibrium entails "separating equilibrium", where those with different characteristics are effectively fully identified by the actions/choices they make, more generally, the equilibrium, when it exists, may entail pooling or partial pooling, when there is costly verification or in multiple period models where the uninformed party cannot commit not to use information gleaned in one period subsequently; (c) the equilibrium is not in general constrained Pareto efficient; (d) the precise specification of the "game" describing the interaction between the two parties matters;[43] and (e) there exists a variety of interventions by government that can be welfare enhancing – price interventions and disclosure and fraud laws.

4.   HOW MARKET DEVELOPMENTS ARE UNDERMINING DECENTRALIZATION AND INCREASING INFORMATION BURDENS

As we have seen, a central idea of the "information" revolution in economics is that prices do not convey all the relevant information. Changes in market structure can accordingly change the informational efficiency of the economy by increasing or decreasing the burden put on various sources of information and can affect the effectiveness of these sources in communicating the relevant information. There can be a tension here. Seeming improvements in markets, *ignoring information*, may actually worsen overall economic performance once the informational burdens, or the consequences of not getting relevant information, are taken into account: still another instance of the general theory of the second best.

Trade liberalization is an obvious example. There are standard arguments for the benefits of trade integration.[44] But as markets get integrated, the relevant set of information for market participants gets greatly increased; because volatility of prices, for example, may increase, market participants will want more information about factors affecting price. Whether welfare is increased can thus be ambiguous.

An even better example is the creation of derivatives, designed to "complete markets", or fill in for *some* of the missing Arrow-Debreu risk markets. But especially the way they have been constructed, they have increased information burdens enormously and undermined effective decentralization, especially the way they have been constructed. One of the great

virtues of the market economy was supposed to be that no one had to know the preferences or technology of others. Prices were, as we have noted, sufficient statistics, conveying all the relevant information. But if, say, bank A has taken out a credit default swap (CDS) with bank B betting that C will default, A's financial position depends on B's ability to fulfil that contract (the counterparty risk); but if B has CDS's with D and F, then B's ability to fulfil its contract depends on judgements concerning D and F's ability to fulfil their contracts, and that will depend on all of the CDS's that they hold. In short, risk assessment of *any* bank requires knowledge of the balance sheets (including derivatives holdings) of *all* banks and firms that are intertwined in this financial network, almost surely all large banks and many other large enterprises. Decentralization has been undermined; overall information burdens increased enormously.[45]

Making matters even worse is that creating these new financial products opens up new betting opportunities, which at least some firms and individuals will take advantage of. This increases the volatility in their wealth positions, increasing macroeconomic volatility and increasing still further information burdens (Guzman and Stiglitz, 2020, 2021a, 2021b).[46]

5.   GOING BEYOND MARKETS: MECHANISM DESIGN

Markets are one way of allocating resources. Individuals reveal their preferences by their purchases and sales at particular prices; and firms, their technological capacities. With prices set according to certain rules, i.e., to clear markets, the outcomes are efficient. Once it became clear that with imperfect, asymmetric, and endogenous information price signals were no longer enough to coordinate economic activity efficiently, and that asymmetric information is a pervasive feature of economic environments, there was a push to explore other ways for market participants to communicate (e.g., information about preferences or technology) and make inferences, and other ways for allocating resources. This is especially important when there are too few market participants to make the price-taking assumption of competitive equilibrium persuasive, and in particular where there is a single agent – a monopolist or the government – that can create a "mechanism" to efficiently extract relevant information in ways that maximize *its* objective function.

Shortly after the development of the basic information models, tools from mathematics – in particular, game theory – were beginning to be used more widely by economists, with extraordinary results. *Mechanism design* (sometimes known as "implementation theory") can be thought of as the flipside of game theory – instead of starting with a game and looking for a solution, one starts with a desired outcome (the solution) and tries to find games (or, more technically, "game forms" – games where everything but the payoffs are specified) in which a specific desired outcome is an equilibrium. The participating agents send "messages" to a centralized "mechanism", which then, depending on which messages the agents send, decides on an allocation – quantities and prices.[47] The task of the mechanism designer is, given a particular social choice function, to find mechanisms in which the outcome of the game that the agents play, with each maximizing her own welfare according to the rules of the game, maximize this social choice function.

Part of the reason "implementation" theory is called so is that it analyses, given a fixed planner's preference ranking over outcomes, which outcomes can be implemented by an appropriately chosen game, and how – what does the game form look like, what are the neces-

sary transfers, and so on. The constraints that the planner faces may be complex – agents may have private information, may have aligned or misaligned preferences – so the implementation problem is difficult in general. Because of this, this literature is rife with impossibility theorems: Arrow, Gibbard-Satterthwaite, Muller-Satterthwaite, and others. So, the contribution is not always of the form "this is the optimal mechanism", but rather of the form "the exists no mechanism that accomplishes all goals", or that the only possible mechanism has very undesirable properties (such as being "dictatorial" – always ranking outcomes according to one agent's preferences, and always disregarding those of the others).

The literature on mechanism design/implementation theory, like that of game theory more generally, shows that results are highly sensitive to assumptions about information, about the players' preferences, and about players' beliefs, including their beliefs about the beliefs of the other players; in particular, the mechanism designer is often assumed to have a lot of information. Below, we consider some of the most salient aspects.

(a) The design of economic mechanisms relies on agents' beliefs about an uncertain, payoff-relevant random variable. The approach to studying games of incomplete information (games where players are uncertain about the payoffs that other players obtain – i.e., they do not know which game they are playing) is due to Harsanyi (1967, 1968a, 1968b); instead of analysing incomplete information games, he proposed analysing games of complete but imperfect information (where players know the payoffs, but do not observe previous moves of other players). His construction involved augmenting the game with an additional player, called Nature, who moves first, and chooses realizations of a random variable for each player according to a distribution known to all players; this realization is called the player's *type*. A type summarizes all of the information about the player – whether she is of high or low ability, has high or low costs, etc. This also allows a player to form beliefs about others' types, and their beliefs about one's type, and beliefs about beliefs – an infinite *hierarchy of beliefs*; a type is assumed to include one's entire hierarchy of beliefs. The languages of hierarchies of beliefs, and type spaces, are complementary.[48] The profile of types (one for each player) is simply a vector of types. In Harsanyi's own words:

> we can regard the vector c_i as representing certain physical, social, and psychological attributes of player i himself in that it summarizes some crucial parameters of player i's own payoff function U_i as well as the main parameters of his beliefs about his social and physical environment … the rules of the game as such allow any given player i to belong to any one of a number of possible types, corresponding to the alternative values of his attribute vector c_i could take. … Each player is assumed to know his own actual type but to be in general ignorant about the other players' actual types. (Harsanyi, 1967, pp. 171–172)

(b) But as early as 1972 Leo Hurwicz recognized the need for mechanisms that do not depend on assumptions about the agents' characteristics, what he called "nonparametric" mechanisms (Hurwicz, 1972).

(c) Most importantly, mechanism design depends on individuals' beliefs about other agents' beliefs (hence the term "hierarchies" of beliefs). Aumann's (1976) work on "common knowledge" has been crucial for clarifying what it means for agents' beliefs about others (and beliefs about beliefs – second-order beliefs, and beliefs about beliefs about beliefs, and so on) to be commonly known. An event *E* is *common knowledge* between two agents if agent 1 knows it, agent 2 knows, agent 1 knows that 2 knows it, and so on, ad infinitum. Aumann (1976) showed

that if individuals have common priors, and their posteriors are common knowledge, their posteriors must be equal: they cannot agree to disagree.[49]

An approach based on common knowledge has been criticized, most notably by Robert Wilson. Wilson emphasized the importance of not making strong informational assumptions, such as common knowledge, in hopes of generating better, more applicable, theories. He wrote, outlining what has become known as the *Wilson doctrine*, "only by repeated weakening of common knowledge assumptions will the theory approximate reality" (Wilson, 1985, 1987).

An example of work in this spirit is Dasgupta and Maskin (2000), who focus on "detail-free" auction rules "that are independent of the details – such as functional forms or distribution of signals – of any particular application and that work well … in a broad range of circumstances" (Dasgupta and Maskin, 2000, p. 347).

5.1   Blackwell's Theorem: A Cornerstone of Information Measurement

We have been discussing the idea of "information", and using terms like "information structure"; while being more or less "informed" may have a clear and intuitive meaning in many of the early "simple" applications, this is not so in general. There is, in fact, a large theory that provides measuring the amount of information,[50] or at least ascertaining whether there is more information in one situation than in another, in much more complicated settings.

A cornerstone of this theory is David Blackwell's (1951, 1953) work. Observe first that generally speaking, different decision makers will value different pieces of information differently, depending on their preferences, other assets, and risk attitudes. For instance, an investment analyst who is interested in the federal funds rate may not value information about the climate, which may be very valuable to another decision maker. Thus, there can be no unanimous ranking of all information structures, by which everyone could agree that one set of information is more informative or "better" than another. But is there *any* context in which *all* expected utility-maximizing decision makers would agree that one set of information is more informative than another? As Blackwell (1951, 1953) showed, the answer is yes: information structure A is more informative than information structure B if any payoff that is attainable under B is also attainable under A. This "payoff richness" appears to be a natural criterion of evaluating information for an economist.

There is another perspective one can take. Suppose one takes information structure A and adds pure noise to it, to obtain information structure B. This seems like a very plausible requirement for A to be more informative than B (we say that A is "sufficient" for B). The intuition is that one can turn A into B *without knowing anything* about the true state. A striking result of Blackwell is that these two ways of evaluating informativeness of information structures are, in fact, identical. Information structure A is payoff richer than B for any utility function if and only if A is sufficient for B.[51] This result provides not only a completely unambiguous ranking of information – payoff richness – but also links it to a mathematically tractable and statistically appealing way of ranking information – sufficiency.[52]

The strength and power of this theorem – it applies to all expected utility maximizers – is also its shortcoming. The Blackwell order is not only partial (meaning most information structures are not ranked) but, loosely speaking, "very" partial. There have been many attempts to "complete" the Blackwell order (Frankel and Kamenica, 2019; Kosenko, 2022; Cabrales et al., 2013, 2017; Pęski, 2008; Mu et al., 2021); no single completion has been found to be appealing in all contexts.

Using the Blackwell framework, Radner and Stiglitz (1984) were able to derive a striking result about the value of information under very weak assumptions: it is always non-concave. Ignorance may not be bliss, but it is at least a local optimum. Much of standard economics (exemplified by Debreu, 1954, 1959) is based on assumptions of concavity. In a world with endogenous information, it is hard to hold that assumption.[53]

5.2   Auctions

Auctions are perhaps the most widely developed example of a commonly used mechanism for allocating resources, as was recognized by the 2020 Nobel Memorial Prize awarded to Robert Wilson and Paul Milgrom – two of the leading figures in auction design.[54] The literature on mechanism design has called attention to better ways of designing auctions under increasingly complex situations. Many of these ideas have been put into practice, with mixed results. While there have been some successes on real-world auctions (radio spectrum auctions in the US, auctions for 3G licences in the UK in 2000), there have also been many disappointments, with unexpectedly low prices (including the Australian satellite TV auction, the New Zealand radio spectrum auction, and the US Environmental Protection Agency's SO2 auction), some of which can be attributed to collusion by the bidders, or poor auction rules design.

5.3   Matching, or Markets without Prices

Alvin Roth was awarded a Nobel Memorial Prize in economics for his work on what can be viewed as a particular class of mechanisms aimed at "matching", e.g., kidney donors with recipients,[55] scarce medical equipment (such as ventilators) with patients who need them,[56] recent MD graduates with hospitals for internships, or high school applicants with public high schools.[57] Matching theory is useful in situations where there is a supply and a demand of an item or service (or more generally, just two sides of the exchange), but there is no price (for legal, moral, or technical reasons). Some, like the kidney donor programme, have been huge successes; others, like the high school programmes, have been more problematic. Understanding better the nuances that lead to success or failure is a subject of ongoing research.[58]

Even when there is no central authority that can design a relevant mechanism for resource allocation (as in the cases just discussed), mechanism design may provide insights; for instance, the gains to trade that could be achieved under an optimal design provide an upper bound to those achievable.

5.4   Consequences of Information Assumptions for Mechanism Design

We noted earlier the sensitivity of many of the results in mechanism design to the particular informational assumptions employed. The literature has naturally asked how the optimal mechanism might be changed as the information available to the economic agents changes (for example, in agents' beliefs about their opponents' valuations). In particular, Bergemann and Morris (2005) have created a framework of "robust" mechanism design, investigating, for instance, how the design of auctions or of systems of price discrimination depends on the assumptions of common priors and common knowledge.[59] In particular, (the main question) they consider is this: suppose that the social planner (or mechanism designer) has a social

choice function she would like to implement. (A note on the nomenclature: this literature distinguishes between "social welfare functions", the domain of which is individual preferences, and the range of which are *rankings* of alternatives, "social choice functions", which have the same domain, output *single* alternatives (not a ranking, as with social welfare functions), and "social choice correspondences", the domain of which is the same, but the output may be multi-valued – they produce a *set* of unranked acceptable alternatives.) What the planner should choose to maximize her social choice function may depend on variables that are only known (initially) to agents. Obviously, the planner would like the agents to convey this information, but they don't just freely do so. They have to be induced to do so through a "game". Bergemann and Morris (2005) ask: When does there exist a game with the property that when the agents play that game, for any information they may have, each agent has an incentive to tell the truth about her information, if she expects others to tell the truth, *whatever their information turns out to be*, and that the outcome of this game always implements the social choice function, i.e. entails, when all the information is revealed, the planner taking an action consistent with what maximizes her utility given the information (state of the world)? It is the requirement that the agents have an incentive to tell the truth if they expect others to tell the truth, whatever others' information turns out to be, that makes this *robust* (*ex post* implementable) – the agent has no incentive to lie, provided others do not lie, *regardless* of their information. This is a very strong concept (akin to the definition of a dominant strategy – a strategy that is a best response to *any* strategy of one's opponent). It is obviously desirable to find such a robust mechanism, because robustness in this sense means that other outcomes are extremely unlikely to result in practice, by construction.

5.5   Applications of Robust Mechanism Design

Two important theoretical applications of robust mechanism design, where the design reflects the lack of information of the designer, are Carroll (2015) and Guo and Shmaya (2019). Carroll (2015) studies a moral hazard problem where the principal is uncertain about the set of actions available to the agent. "She [the principal] knows some of the available actions, but other unknown actions may also exist and our principal does not even have a prior belief about these unknown actions". He finds that the solution (under very mild assumptions concerning preferences and the ability to impose taxes, e.g., in outcome-dependent ways, under which the first best outcome could be trivially implemented) to this problem is linear contracts – i.e., that if a contract is judged by its worst possible performance, linear contracts uniquely provide the highest guaranteed return to the principal (i.e. it is the "maximin" solution). Carroll (2015) uses this to explain the pervasiveness of linear contracts in the real world – where the principal knows that there may be eventualities that matter that are not even conceivable now, and obviously can't form beliefs about things she does not know exist – as opposed to "optimal" contracts from non-robust contract theory that often take complicated forms.[60]

Guo and Shmaya (2019) study from a non-Bayesian point of view the same problem of monopoly regulation (again viewed as a principal-agent problem) that Baron and Myerson (1982) studied from a Bayesian approach. They identify policies that minimize "regret" of the regulator – the difference between what he could have obtained with perfect information and what he actually obtains. The result is an interesting combination of price caps (to benefit the consumer) and piece-rate subsidies (to incentivize the producer); the precise amount and

combination of these tools depends on how much weight the regulator puts on the two sides of the market.

6.   CONCLUDING REMARKS

The analysis of environments with costly information has provided enormous insights into the workings of virtually all markets and of the economic system as a whole. It has provided explanations of phenomena which could not be explained by the standard theory assuming perfect information. We have seen the multiple ways in which early results concerning information revelation, screening, and signalling have been reinforced and modified, with results in some cases being shown to be more robust than was at first thought to be the case, while in others, more fragile. Most importantly, the conclusion that markets are, in general, not (constrained Pareto) efficient has, if anything, been strengthened, and indeed, we have seen how many of the reforms of recent years may have actually decreased welfare, once one takes into account the consequences of the imperfections and endogeneity of information. Recent work has highlighted the potential of information asymmetries to magnify distributional differences, entrench market power, and make efficient regulation more difficult, but has also suggested ways of at least partially overcoming these problems and improving resource allocations. In particular, we have seen how a greater understanding of how the problems posed by information asymmetries can be overcome continues to generate new insights; ingenious new mechanisms have emerged that implement desirable outcomes when intuition might suggest such outcomes may be impossible.

The companion chapter that follows explores a particularly fruitful strand of work – where information is conveyed not just by making inferences based on choices and actions, but also by direct communication.

ACKNOWLEDGEMENTS

NOTES

1.   See, e.g., Stiglitz (1985b) and Stiglitz and Wolfson (1988). There is a wealth of other applications, most which we cannot explore in this short chapter. These include regulation (itself a huge literature, including Sappington and Stiglitz, 1987a), the exercise of and regulation of market power (e.g., Baron and Myerson, 1982; Rey and Tirole, 1986; Laffont and Tirole, 1986, 1988, 1990), and privatization (Sappington and Stiglitz, 1987b).
2.   Indeed, there have been multiple surveys, both of the field in general (e.g., Stiglitz, 1975a, 2000, 2002, 2013b, 2020; the introductory essays in Stiglitz, 2009, 2013a; and Kamenica, 2017) and its application to particular subdisciplines, including labour (see Ashenfelter and Card, 2011a, 2011b), product (see Stiglitz, 1989), and insurance (see Dionne, 2013) markets. Veldkamp (2012) discusses

applications of imperfect information in macroeconomics and financial markets. By the same token, we necessarily must be selective in the references we choose to cite. We make no attempt either at completeness or comprehensiveness.

3. Although, notably, this insight fails to hold with more than two decision makers, as the following discussion will make clear.

4. The two key properties of a pure Samuelsonian public good are non-rivalrous consumption (what one person consumes doesn't subtract from what another can) and non-excludability. Information has both properties, but we have noted even if it were possible to exclude partially (as patents do for certain kinds of information), the resulting market equilibrium is not efficient.

5. This chapter notes the inefficiency associated with obtaining a particular kind of information, knowledge associated with increased efficiency in production, namely, research and development. Thus, the market failures noted here in the context of screening models are also relevant for investments in innovation. These inefficiencies are more extensively discussed in Stiglitz and Greenwald (2015).

6. Akerlof (1970) showed that with information asymmetries, there might not exist a market for used cars, a specific application of a more generalized "no trade theorem" (Milgrom and Stokey, 1982; Stiglitz, 1982b). For applications of these ideas to labour markets, see Greenwald (1979, 1986). For applications to equity markets, see Greenwald et al. (1984). There are also fixed costs associated with establishing and running markets, providing an alternative explanation for the absence of markets.

7. Diamond (1967) provided a weaker set of sufficient conditions for a weaker notion of "constrained Pareto efficiency", but Stiglitz (1982a) showed that whenever there were more than one good or bankruptcy costs, markets were not constrained Pareto efficient even in the weaker sense that Diamond had proposed. Earlier, Stiglitz (1972) had shown even in the more restrictive world of mean-variance (the basis of the capital asset pricing model) the economy was not efficient. More generally, Greenwald and Stiglitz showed that *given the set of markets in existence*, there exist interventions in the market allocation of goods that could improve the welfare of some without decreasing that of others. Their result can be seen as a generalization not only of Stiglitz (1982a), but also of Newbery and Stiglitz (1982), who show the (constrained) inefficiency of markets even with rational expectations.

8. Newbery and Stiglitz (1984) illustrate showing that opening up new opportunities to trade between two economies may lower the welfare of all individuals in both economies if there are not risk markets. Later, we expand on the specific reasons that adding additional securities – short of a complete set of risk markets – may be welfare decreasing.

9. Or when there are collateral constraints, as in many recent macroeconomic models.

10. The economics of information help explain the widespread institutional arrangement in agriculture of sharecropping, which from the perspective of standard economics seemed a peculiar economic arrangement, significantly reducing workers' incentives. Indeed, sharecropping provided the paradigm of the "agency" problem, where effort was unobservable, but workers were risk averse (Stiglitz, 1974a). Braverman and Stiglitz (1982) showed in that context that there would be interlinking of markets – decentralization failed.

11. For an application of these externalities to the Covid-19 pandemic, see Guzman and Stiglitz (2021c).

12. There is a vast literature claiming that markets are informationally efficient (see, e.g., Fama, 1970, 1991), but an even more important literature establishing that they are not (Shiller, 1981, 2000). Grossman (1976, 1977) and Grossman and Stiglitz (1980) not only established that markets could not be informationally efficient in transmitting information from the informed to the uninformed, but that they even failed to aggregate disperse information well. Similarly, Gale and Stiglitz (2013) showed that futures markets are almost always informationally inefficient. One of the most curious Nobel Prize awards was that of 2013, awarded both to Fama, for his work on informationally efficient markets, and to Shiller, for showing that markets were not informationally efficient.

13. First noted in Stiglitz (1969, 1974c, 1982d). Since then there has been an enormous literature on the efficiency wage model (see, e.g., Weiss, 1980) and on credit rationing (Stiglitz and Weiss, 1981). For an overview, see Stiglitz (1987).

14. For instance, many of the seeming anomalies in macroeconomics arise because market prices (say, interest rates) and shadow prices may move in different directions in the presence of credit rationing.

15. See, for instance Diamond (1971) who shows that even with small search costs, the market price will be the monopoly price. Stiglitz (1985a) shows analogous results in labour markets, and established that there may be multiple equilibria wage distributions. Stiglitz (2013a) shows, moreover, that in the Diamond model, in the absence of heterogeneity, there is in general no equilibrium.

16. Walras' Law held that if there are *N* markets, and *N-1* clear, the *N*th must clear.

17. While Rothschild and Stiglitz (1976) showed that in their one-period model based on self-selection and contract exclusivity, there could not be a pooling equilibrium, in other contexts, there can be a pooling equilibrium. See the discussion below. Note that in Stiglitz (1975a) the market equilibrium *may* be inefficient even with a single commodity; in the RS (1976) analysis, the market equilibrium, when it exists, will be Pareto efficient, but only with a *single* commodity. With multiple goods, self-selection equilibria are generally inefficient (see Greenwald and Stiglitz, 1986, and Arnott et al., 1994).

18. In the simple early models, such as Spence (1973), Stiglitz (1975a), and Mirrlees (1971), more able individuals were proportionately better at *every* task; there was no comparative advantage. But when there exists comparative advantage, with one type of individual's having better *relative* performance at some task than other types, but poorer relative performance in others, knowing individuals' abilities allows a better allocation of individuals to different tasks. There is a social return to having information about individuals' relative abilities. Of course, there will always be some comparative advantage, except in the trivial case – which is what the literature focused upon. The results of that literature appear robust, so long as the magnitudes of differences in comparative advantages are not large. Surprisingly, little of the literature has explored such situations.

19. Assuming there is only one good. As we have noted, if there is more than one commodity, the Greenwald-Stiglitz theorem applies, so that the economy is not Pareto efficient (see Greenwald and Stiglitz, 1986).

20. Throughout the chapter, the exposition focuses on identifying differences in individual abilities; the discussion could have been exposited as well in terms of differences in project returns, risks of individuals seeking insurance, or qualities of products.

21. This might not be a problem if one could buy insurance *before* the information were available to any party, e.g., before the individual is born. But that is not possible – partly because genetic information about parents (or asymmetries about such information) may lead to the unravelling/non-existence of that market.

22. Similarly, Stiglitz (1985a) and Salop and Stiglitz (1977, 1982) show that the market may create price dispersions, even when there is no intrinsic difference among firms. See also Stiglitz (2013a). When it does this, it imposes unnecessary search costs on consumers. See the discussion below in sections 3.11 and 3.12.

23. See, *inter alia*, Thaler and Benartzi (2004), Benartzi and Thaler (2007), and DellaVigna and Malmendier (2006).

24. There is an analogy here to the costs imposed on consumers in markets with endogenous price dispersion noted in note 22 and discussed further below. The high price firm exploits the high search cost individuals, but simultaneously imposes a search cost externality on the low search cost individuals. There is a further rationale for compulsory disclosure when verification is costly. We showed that the pooling equilibrium is Pareto superior to the screening equilibrium *if the supply of the two types is fixed*. But the pooling equilibrium will lead to an oversupply of "bad" products and an undersupply of good products, since the price they receive is the same, i.e., *incentives* are distorted. Taking these into account, the pooling equilibrium may well be welfare-inferior; and disclosure mandates may be necessary and desirable.

25. Obviously, there are some circumstances in which the seller may not be fully informed about the characteristics of his or her product; but typically, it is more efficient/less costly for the seller to gather relevant information than for a multitude of buyers. On the other hand, the information may be less trustworthy, especially in the presence of imperfect and costly verifiability (so mild deviations from the truth are hard to prosecute). In such cases, third-party provision of information (like Consumers' Reports) may be desirable, either as a substitute or complement. Thus, the testing of

drugs for safety and efficacy by the pharmaceutical companies selling those drugs has been questioned, with some advocating that third-party testing should play a more important role (Jayadev and Stiglitz, 2008, 2010).

26. Thus, securities laws requiring disclosure of risks have proven to be less effective than hoped, because firms bury the real and important risks to which attention should be drawn within a long list of ordinary risks to which investors are exposed.

27. On the role of guarantees, see Grossman (1981), and Heal (1977). A money back guarantee by itself does not suffice in the presence of costly enforcement. Note that increasing the costs of enforcement (e.g., by not allowing class action suits) increases opportunities for low-information ("bad") equilibrium to occur (e.g., where firms *exploit* poorly informed consumers) – and increases the likelihood of a no-trade equilibrium (with critical markets being absent).

28. Or in the case of confinement, where the length of confinement is sufficient to deter fraudulent behaviour.

29. Stiglitz and Weiss (1994) further clarify the distinction between signalling and screening models. In Stiglitz (1977), a single firm (a monopolist) constructs the choice set to enable the differentiation among individuals of different types. RS considers the more complicated situation where the choice set emerges out of a competitive equilibrium.

30. Similarly, earlier we noted that a guarantee is a costly signal. The cost to providing a guarantee for the good product, one where the probability of a critical defect is lower, for example, is less than for the bad; but, as we saw in the previous section, the guarantee can be thought of as part of the "verification of truth". To put it rather inelegantly, the party providing the guarantee is putting his money where his mouth is.

31. As we noted earlier, a key insight in these screening and signalling models is that the social costs and benefits of signalling and screening may differ markedly from the private costs and benefits.

32. In the Spence model, equilibrium is defined simply as a set of self-confirming beliefs: given wage differences between the educated and the uneducated, those choosing to get educated (not educated) had productivities precisely corresponding to those beliefs. There might, of course be multiple such sets of equilibrium beliefs. These refinements eliminate the multiplicity of equilibrium, doing so typically by asking what inferences the uninformed would make were they to see an action that was deviant. See also Stiglitz (2002), putting forward the almost obvious point: in any putative signalling equilibrium, the least able individual has nothing to lose by simply investing nothing in education. The *worst* that could happen is that he would be treated as if he were the least able – and then he would still be better off than he is in the signalling equilibrium where he invests in education. So too, the more able could invest far more in education than needed to "signal" that he is of better quality. That, in Spence's sense, would be an equilibrium. But it is obvious that any high ability individual who invested at least enough such that no low individual would invest that much (given their cost differences and putative returns) would do so, would still convey (signal) that he is high ability.

33. The Riley "equilibrium" is a reactive equilibrium, inconsistent with the spirit of competitive analysis that motivated, for instance, Rothschild and Stiglitz. Acevedo and Gottlieb (2019) explore the non-existence of the Riley equilibrium.

34. But this does not address the issue of whether full revelation is efficient. In general it is not. See Stiglitz (2009). This helps explain why there does not exist a screening equilibrium with a continuum of individuals. Recall the discussion above where, with two types, an equilibrium only exists if the types are different enough. Obviously, with a continuum, individuals are arbitrarily similar to others. It can be shown that it does not pay those near the least able individuals to distinguish themselves from that type.

35. For an overview of the repeated games literature, see Mailath and Samuelson (2006). Casella (2005) and Jackson and Sonnenschein (2007) discuss other mechanisms. The titles of the latter two works are rather evocative: "Storable Votes" and "Overcoming Incentive Constraints by Linking Decisions".

36. As in Shapiro and Stiglitz (1984) or Stiglitz and Weiss (1983). These results are related to the earlier discussion of decentralization: here, intertemporal linkages are critical. See also Shapiro (1982) and Klein and Leffler (1981). Under certain conditions in labour markets, firms can punish workers for

whom there has been a (noisy) signal of shirking by lowering wages (rather than terminating the contract). See Rey and Stiglitz (2023) and Arnott et al. (1988).

37. In Chapter 3, we consider models with direct communication, giving rise to the problem of "information design". An information design approach to the sequential screening problem is studied by Krämer and Strausz (2015) and Heumann (2020).

38. This is analogous to "robust" mechanism design discussed below.

39. In Akerlof's model, where individuals just purchase or sell one car, quantities are not relevant. In the insurance market (and many other markets), there may be other aspects of the economic transaction that are observable and convey information.

40. The literature on equilibrium with non-exclusivity is large and complex. See Arnott and Stiglitz (2013a, 2013b) and the discussion in KSY (2023).

41. In both cases, for instance, the maximization problem entails self-selection constraints. There are many institutional details that need to be incorporated, and when this is done, there are salient differences, accounting for the huge literatures which have developed in each of these separate fields. Early contributions noting the formal similarity include Sappington and Stiglitz (1987a) and Stiglitz (1982c).

42. Rothschild and Stiglitz (1976). In fact, when the types are not too different – as when there is a continuum of types – there may never exist an equilibrium. See Riley (1975, 1979, 2001) and Stiglitz (2009). Dasgupta and Maskin (1986a, 1986b) provide a game theoretic formulation in which there is a mixed strategy equilibrium.

43. For instance, it matters whether the informed party "moves" before the uninformed (as in the typical signalling game) or vice versa (Stiglitz and Weiss, 1994).

44. As we noted earlier, those do not hold in the absence of perfect risk markets. Indeed, trade liberalization can be Pareto inferior (Newbery and Stiglitz, 1984). The discussion here goes beyond Newbery and Stiglitz in noting that trade integration leads to (greater) price variability, and therefore the value of information about what prices will be increases.

45. Indeed, derivatives may greatly increase the complexity of markets and give rise to a fundamental indeterminacy of equilibria. See Roukny et al. (2018) and Battiston et al. (2016).

46. These are not the only examples of misguided market "reforms". In some quarters, there is support for moving from bank lending to capital markets, in the belief that the latter do a better job in spreading risks. But as Grossman and Stiglitz (1980) and Stiglitz and Greenwald (2003) emphasize, because of informational spillovers, capital markets won't have sufficient incentives for gathering information. See also Stiglitz (1992).

47. There are various strands within this literature where the mechanism designer may not be able to completely control the behaviour of agents (the allocations) (as in our earlier discussion of the insurance market where customers may purchase secret insurance); and where the mechanism designer may not be able to commit to a particular rule before the messages are sent. See, *inter alia*, Rahman (2012), whose work is evocatively titled "But Who Will Monitor the Monitor?"

48. Mertens and Zamir (1985) and Brandenburger and Dekel (1993) provided a critical mathematical foundation for these ideas, by constructing a *universal type space*, such that any "reasonably consistent" infinite hierarchy of beliefs can be represented by a type in this universal type space, and along the way explicitly constructing a mathematical mapping that maps types to infinite hierarchies of beliefs, and vice versa. This mapping is one-to-one (injective), onto (surjective), continuous, and has a continuous inverse (a "homeomorphism"); by virtue of satisfying all of these conditions, this map "preserves" all the features of one space when mapping it to another, and vice versa.

49. Of course, without common priors, even sharing information will not lead to common posteriors. One can then agree to disagree, as was recognized and discussed by Aumann (1976). Chapter 3 discusses some of the problems of societal polarization which arise in the presence of different priors. The assumption of common knowledge is standard in the rational expectations macroeconomic literature and has provided a central point of critique of dynamic stochastic general equilibrium (DSGE) models (see, e.g., Stiglitz, 2018 and Guzman and Stiglitz, 2021a, 2021b).

50. These ideas are related to, but different from "information theory" where the focus is on the information transmitted through a "channel" – an idealized communication device.

51. In more mathematical terms, there exists a "garbling" matrix (in the finite case), or a Markov kernel (in the infinite case) that when applied to A, yields B.

52. The link between Rothschild and Stiglitz's work (1970, 1971) ranking probability distributions *for any risk averse* individual and that of Blackwell should be clear; they provide additional equivalent characterizations, as does Atkinson (1970), in the context of rankings of distributions of income. For extensions in that context, beyond separable utility functions, see Rothschild and Stiglitz (1973). For analyses of behavioural implications, see Rothschild and Stiglitz (1971) and Diamond and Stiglitz (1974). We make one clarifying remark – while Rothschild and Stiglitz (1970, 1971) characterize mean-preserving spreads of distributions of monetary lotteries, Blackwell informativeness is equivalent to mean-preserving spreads of distributions of beliefs; in the context considered by Rothschild and Stiglitz, the two are equivalent.
53. Similar results were shown to hold in models with moral hazard (Arnott and Stiglitz, 1988, 2013a, 2013b).
54. Even before that, the Nobel Memorial Prize was awarded to William Vickrey (along with James Mirrlees) in 1996, partly for his work on second price auctions (Vickrey, 1961).
55. See Roth (2003) and Roth et al. (2004, 2005a, 2005b, 2006, 2007, 2008).
56. See Pathak et al. (2020) for an example motivated by the coronavirus pandemic; different US states had different priority rules for assigning scarce medical resources, which aimed to balance various ethical and practical considerations. Pathak and co-authors proposed a better mechanism, which was subsequently adopted by the University of Pittsburgh Medical Center.
57. Roth and Sotomayor (1990) summarize much of the theoretical work in this area; Teytelboym et al. (2021) provide a more recent summary.
58. See also the references to work by Atila Abdulkadiroglu, Parag A. Pathak, Alvin E. Roth, Tayfun Sonmez and M. Utku Ünver.
59. The "robustness" of mechanism design has a similar flavour to the belief-free equilibria of the context of repeated games, where the term "robust" was also used, but the precise meaning differs because the contexts are different. Belief-free equilibria are robust to individuals' beliefs about others in repeated games; robust implementation occurs when a mechanism designer implements a social choice function as an equilibrium where individuals tell the truth, provided they expect others do as well, for any information they may have.
60. See also Diamond (1998), who also finds linear contracts to be optimal in a different, simpler, though more restrictive setting. See also Allen (1985), who focuses on the implications of non-observable trades, and which may be a better explanation of linear contracts than that provide by Carroll's theory, since the critical unknown, the weather, has (particularly before the onset of climate change) a relatively well-defined probability distribution.

REFERENCES

Acevedo, E. & Gottlieb, D. (2019). An Example of Non-existence of the Riley Equilibrium in Markets with Adverse Selection. *Games and Economic Behavior*, 116(C), 152–157.

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

Akerlof, G. A. (1976). The Economics of Caste and of the Rat Race and Other Woeful Tales. *The Quarterly Journal of Economics*, 90(4), 599–617.

Allen, F. (1985). On the Fixed Nature of Sharecropping Contracts. *Economic Journal*, 95(377), 30–48.

Arnott, R., Greenwald, B., & Stiglitz, J. E. (1994). Information and Economic Efficiency. *Information Economics and Policy*, 6(1), 77–82.

Arnott, R., Hosios, A. J., & Stiglitz, J. E. (1988). Implicit Contracts, Labor Mobility and Unemployment. *American Economic Review*, 78(5), 1046–1066.

Arnott, R. & Stiglitz, J. E. (1988). The Basic Analytics of Moral Hazard. *Scandinavian Journal of Economics*, 90(3), 383–413. Reprinted in *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*. Oxford: Oxford University Press, 2009, pp. 417–443.

Arnott, R. & Stiglitz, J. E. (1991). Moral Hazard and Non-Market Institutions: Dysfunctional Crowding Out or Peer Monitoring. *American Economic Review*, 81(1), 179–190. Reprinted in *Selected Works of*

Joseph E. Stiglitz, Volume I: Information and Economic Analysis. Oxford: Oxford University Press, 2009, pp. 589–603.

Arnott, R. & Stiglitz, J. E. (2013a). Equilibrium in Competitive Insurance Markets with Moral Hazard. In *Selected Works of Joseph E. Stiglitz, Volume II: Information and Economic Analysis: Applications to Capital, Labor, and Product Markets*. Oxford: Oxford University Press, 2013, pp. 660–689. Originally Princeton University Discussion Paper 4, 1987. Also NBER Working Paper 3588, 1991.

Arnott, R. & Stiglitz, J. E. (2013b). Price Equilibrium, Efficiency, and Decentralizability in Insurance Markets. In *Selected Works of Joseph E. Stiglitz, Volume II: Information and Economic Analysis: Applications to Capital, Labor, and Product Markets*. Oxford: Oxford University Press, 2013, pp. 632–659. Originally NBER Working Paper 3642, 1991.

Arrow, K. J. (1964). The Role of Securities in the Optimal Allocation of Risk-Bearing. *The Review of Economic Studies*, 31(2), 91–96.

Arrow, K. J. (1973). The Theory of Discrimination. In O. Ashenfelter and A. Rees (Eds.), *Discrimination in Labor Markets*. Princeton, NJ: Princeton University Press, pp. 3–33.

Ashenfelter, O. & Card, D. (Eds.) (2011a). *Handbook of Labor Economics Volume 4 Part A*. North Holland: Elsevier.

Ashenfelter, O. & Card, D. (Eds.) (2011b). *Handbook of Labor Economics Volume 4 Part B*. North Holland: Elsevier.

Atkinson, A. B. (1970). On the Measurement of Inequality. *Journal of Economic Theory*, 2(3), 244–263.

Aumann, R. (1976). Agreeing to Disagree. *The Annals of Statistics*, 4(6), 1236–1239.

Banks, J. S. & Sobel, J. (1987). Equilibrium Selection in Signalling Games. *Econometrica*, 55(3), 647–661.

Baron, D. P. & Myerson, R. B. (1982). Regulating a Monopolist with Unknown Costs. *Econometrica*, 50(4), 911–930.

Battiston, S., Caldarelli, G., May, R. M., Roukny, T., & Stiglitz, J. E. (2016). The Price of Complexity in Financial Networks. *Proceedings of the National Academy of Sciences of the United States of America*, 113(36), 10031–10036.

Benartzi, S. & Thaler, R. (2007). Heuristics and Biases in Retirement Savings Behavior. *Journal of Economic Perspectives*, 21(3), 81–104.

Bergemann, D. & Morris, S. (2005). Robust Mechanism Design. *Econometrica*, 73, 1771–1813.

Bernanke, B. & Gertler, M. (1989). Agency Costs, Net Worth, and Business Fluctuations. *American Economic Review*, 79(1), 14–31.

Blackwell, D. (1951). Comparison of Experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 2, 93–103.

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

Brandenburger, A. & Dekel, E. (1993). Hierarchies of Beliefs and Common Knowledge. *Journal of Economic Theory*, 59(1), 189–198.

Braverman, A. & Stiglitz, J. E. (1982). Sharecropping and the Interlinking of Agrarian Markets. *American Economic Review*, 72(4), 695–715.

Cabrales, A., Gossner, O., & Serrano, R. (2013). Entropy and the Value of Information for Investors. *American Economic Review*, 103(1), 360–377.

Cabrales, A., Gossner, O., & Serrano, R. (2017). A Normalized Value for Information Purchases. *Journal of Economic Theory*, 170, 266–288.

Calabresi, G. (1970). *The Cost of Accidents: A Legal and Economic Analysis*. New Haven: Yale University Press.

Caplin, A. & Dean, M. (2015). Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review*, 105(7), 2183–2203.

Caplin, A., Dean, M., & Leahy, J. (2022). Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy. *Journal of Political Economy*, 130(6), 1676–1715.

Carroll, G. (2015). Robustness and Linear Contracts. *American Economic Review*, 105(2), 536–563.

Casella, A. (2005). Storable Votes. *Games and Economic Behavior*, 51(2), 391–419.

Cho, I.-K. & Kreps, D. M. (1987). Signalling Games and Stable Equilibria. *The Quarterly Journal of Economics*, 102(2), 179–221.

Courty, P. & Hao, L. (2000). Sequential Screening. Review of Economic Studies, 67(4), 697–717.

Dasgupta, P. & Maskin, E. (1986a). The Existence of Equilibrium in Discontinuous Economic Games, Part I (Theory). *Review of Economic Studies*, 53(1), 1–26.

Dasgupta, P. & Maskin, E. (1986b). The Existence of Equilibrium in Discontinuous Economic Games, Part II (Applications). *Review of Economic Studies*, 53(1), 27–44.

Dasgupta, P. & Maskin, E. (2000). Efficient Auctions. *The Quarterly Journal of Economics*, 115(2), 341–388.

Davila, E. & Korinek, A. (2017). Pecuniary Externalities in Economies with Financial Frictions. *The Review of Economic Studies*, 85(1), 1–44.

Debreu, G. (1954). Valuation Equilibrium and Pareto Optimum. *Proceedings of the National Academy of Sciences of the United States of America*, 40(7), 588–592.

Debreu, G. (1959). *The Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. New Haven: Yale University Press.

DellaVigna, S. & Malmendier, U. (2006). Paying Not to Go to the Gym. *American Economic Review*, 96(3), 694–719.

Delli Gatti, D., Gallegati, M., Greenwald, B., Russo, A., & Stiglitz, J. E. (2012). Mobility Constraints, Productivity Trends, and Extended Crises. *Journal of Economic Behavior & Organization*, 83(3), 375–393.

Diamond, P. (1967). Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment. *Journal of Political Economy*, 75(5), 765–766.

Diamond, P. (1971). A Model of Price Adjustment. *Journal of Economic Theory*, 3(2), 156–168.

Diamond, P. (1998). Managerial Incentives: On the Near Linearity of Optimal Compensation. *Journal of Political Economy*, 106(5), 931–957.

Diamond, P. & Stiglitz, J. E. (1974). Increases in Risk and in Risk Aversion. *Journal of Economic Theory*, 8(3), 337–360.

Dionne, G. (Ed.) (2013). *Handbook of Insurance*. Cham: Springer.

Edlin, A. S. & Stiglitz, J. E. (1995). Discouraging Rivals: Managerial Rent-Seeking and Economic Inefficiencies. *American Economic Review*, 85(5), 1301–1312.

Ely, J. C., Hörner, J. & Olszewski, W. (2005). Belief-Free Equilibria in Repeated Games. *Econometrica*, 73(2), 377–415.

Emran, M. S. & Stiglitz, J. E. (2009). *Financial Liberalization, Financial Restraint, and Entrepreneurial Development* (SSRN Scholarly Paper No. 1332399). Social Science Research Network. https://doi.org/10.2139/ssrn.1332399.

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2), 383–417.

Fama, E. F. (1991), Efficient Capital Markets: II. *The Journal of Finance*, 46, 1575–1617.

Frankel, A. & Kamenica, E. (2019). Quantifying Information and Uncertainty. *American Economic Review*, 109(10), 3650–3680.

Gale, I. & Stiglitz, J. E. (1989). The Informational Content of Initial Public Offerings. *The Journal of Finance*, 44(2), 469–478.

Gale, I. & Stiglitz, J. E. (2013). A Simple Proof that Futures Markets are Almost Always Informationally Imperfect. In *Selected Works of Joseph E. Stiglitz, Volume II: Information and Economic Analysis: Applications to Capital, Labor, and Product Markets*. Oxford: Oxford University Press, pp. 122–132. (Revised version of "Futures Markets Are Almost Always Informationally Inefficient". Princeton University Financial Research Center Memorandum No. 57, February 1985.)

Geanakoplos, J. D. & Polemarchakis, H. M. (1986). Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete. In W. P. Heller, R. M. Starr, & D. A. Starrett (Eds.), *Essays in Honor of Kenneth J. Arrow, Volume 3: Uncertainty, Information, and Communication*. New York: Cambridge University Press, pp. 65–96.

Greenwald, B. (1979). *The Labor Market as a Market for Lemons*. New York: Garland Publishing.

Greenwald, B. (1986). Adverse Selection in the Labour Market. *Review of Economic Studies*, 53(3), 325–347.

Greenwald, B. & Stiglitz, J. E. (1986). Externalities in Economies with Imperfect Information and Incomplete Markets. *The Quarterly Journal of Economics*, 101(2), 229–264.

Greenwald, B. & Stiglitz, J. E. (1988). Pareto Inefficiency of Market Economies: Search and Efficiency Wage Models. *American Economic Review*, 78(2), 351–355.

Greenwald, B. & Stiglitz, J. E. (1992). Information, Finance and Markets: The Architecture of Allocative Mechanisms. *Industrial and Corporate Change*, 1(1), 37–63.

Greenwald, B. & Stiglitz, J. E. (1993). Financial Market Imperfections and Business Cycles. *The Quarterly Journal of Economics*, 108(1), 77–114.

Greenwald, B., Stiglitz, J. E., & Weiss, A (1984). Informational Imperfections in the Capital Market and Macroeconomic Fluctuations. *American Economic Review*, 74(2), 194–199.

Grossman, S. (1976). On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information. *Journal of Finance*, 31(2), 573–585.

Grossman, S. (1977). The Existence of Future Markets, Noisy Rational Expectations and Informational Externalities. *Review of Economic Studies*, 44(3), 431–449.

Grossman, S. (1981). The Informational Role of Warranties and Private Disclosure about Product Quality. *Journal of Law and Economics*, 24(3), 461–483.

Grossman, S. & Stiglitz, J. E. (1980). On the Impossibility of Informationally Efficient Markets. *American Economic Review*, 70(3), 393–408.

Guo, Y. & Shmaya, E. (2019). Robust Monopoly Regulation. Working Paper. https://arxiv.org/pdf/1910.04260.pdf.

Guzman, M. & Stiglitz, J. E. (2020). Towards a Dynamic Disequilibrium Theory with Randomness. *Oxford Review of Economic Policy*, 36(3), 621–674.

Guzman, M. & Stiglitz, J. E. (2021a). Pseudo-Wealth and Consumption Fluctuations. *The Economic Journal*, 131(633), 372–391.

Guzman, M. & Stiglitz, J. E. (2021b). Economic Fluctuations and Pseudo-Wealth. *Industrial and Corporate Change*, 30(2), 297–315.

Guzman, M. & Stiglitz, J. E. (2021c). The Pandemic Economic Crisis, Precautionary Behavior, and Mobility Constraints: An Application of the Dynamic Disequilibrium Model with Randomness. *Industrial and Corporate Change*, 30(2), 467–497.

Harsanyi, J. C. (1967). Games with Incomplete Information Played by "Bayesian" Players, I–III Part I. The Basic Model. *Management Science*, 14(3), 159–182.

Harsanyi, J. C. (1968a). Games with Incomplete Information Played by "Bayesian" Players, I–III. Part II. Bayesian Equilibrium Points. *Management Science*, 14(5), 320–334.

Harsanyi, J. C. (1968b). Games with Incomplete Information Played by "Bayesian" Players, I–III. Part III. The Basic Probability Distribution of the Game. *Management Science*, 14(7), 486–502.

Hart, S. (2005). An Interview with Robert Aumann. *Macroeconomic Dynamics*, 9(5), 683–740.

Heal, G. (1977). Guarantees and Risk-Sharing. *The Review of Economic Studies*, 44(3), 549–560.

Heumann, T. (2020). Information Design and Sequential Screening with Ex Post Participation Constraint. *Theoretical Economics*, 15, 319–359.

Hirshleifer, J. (1971). The Private and Social Value of Information and the Reward to Inventive Activity. *American Economic Review*, 61(4), 561–574.

Hörner, J. & Lovo, S. (2009). Belief-Free Equilibria in Games with Incomplete Information. *Econometrica*, 77(2), 453–487.

Hörner, J., Lovo, S., & Tomala, T. (2018). Belief-Free Price Formation. *Journal of Financial Economics*, 127(2), 342–365.

Hurwicz, L. (1972). On Informationally Decentralized Systems. In C. McGuire & R. Radner (Eds.), *Decision and Organization: A Volume in Honor of Jacob Marschak*. Amsterdam: North-Holland, pp. 297–336.

Jackson, M. O. & Sonnenschein, H. F. (2007). Overcoming Incentive Constraints by Linking Decisions. *Econometrica*, 75(1), 241–257.

Jayadev, A. & Stiglitz, J. E. (2008). Two Ideas to Increase Innovation and Reduce Pharmaceutical Costs and Prices. *Health Affairs*, 27(Supplement 1), w165–w168.

Jayadev, A. & Stiglitz, J. E. (2010). Medicine for Tomorrow: Some Alternative Proposals to Promote Socially Beneficial Research and Development in Pharmaceuticals. *Journal of Generic Medicines*, 7(3), 217–226.

Jeanne, O. & Korinek, A. (2010). Excessive Volatility in Capital Flows: A Pigouvian Taxation Approach. *American Economic Review*, 100(2), 403–407.

Kamenica, E. (2017). Information Economics. *Journal of Political Economy*, 125, 1885–1890.

Kandori, M. & Obara, I. (2006). Efficiency in Repeated Games Revisited: The Role of Private Strategies. *Econometrica*, 74(2), 499–519.

Klein, B. & Leffler, K. B. (1981). The Role of Market Forces in Assuring Contractual Performance. *Journal of Political Economy*, 89(4), 615–641.

Kohlberg, E. & Mertens, J.-F. (1986). On the Strategic Stability of Equilibria. *Econometrica*, 54(5), 1003–1037.

Korinek, A. (2012). Capital Flows, Crises, and Externalities. In F. Allen et al. (Eds.), *The Global Macro Economy and Finance: Proceedings of the 16th IEA World Congress, vol. 3*. Basingstoke: Palgrave Macmillan, Chapter 5.

Kosenko, A. (2022). Algebraic Properties of Blackwell's Order and a Cardinal Measure of Informativeness. Working Paper.

Kosenko, A., Stiglitz, J. E., & Yun, J. (2023). Bilateral Information Disclosure in Adverse Selection Markets with Nonexclusive Competition. *Journal of Economic Behavior & Organization*, 205, 144–168.

Krähmer, D. & Strausz, R. (2015). Optimal Sales Contracts with Withdrawal Rights. *Review of Economic Studies*, 82, 762–790.

Laffont, J.-J. & Tirole, J. (1986). Using Cost Observation to Regulate Firms. *Journal of Political Economy*, 94(3), 614–641.

Laffont, J.-J. & Tirole, J. (1988). The Dynamics of Incentive Contracts. *Econometrica*, 56(5), 1153–1175.

Laffont, J.-J. & Tirole, J. (1990). Adverse Selection and Renegotiation in Procurement. *The Review of Economic Studies*, 57(4), 597–625.

Leitzinger, J. & Stiglitz, J. E. (1984). Information Externalities in Oil and Gas Leasing. *Contemporary Economic Policy Issues*, 5, 44–57.

Maćkowiak, B., Matejka, F., & Wiederholt, M. (2019). Rational Inattention: A Review. Working Paper.

Mailath, G. & Samuelson, L. (2006). *Repeated Games and Reputations: Long-Run Relationships*. New York: Oxford University Press.

Matějka, F. & McKay, A. (2015). Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review*, 105(1), 272–298.

Mertens, J.-F. & Zamir, S. (1985). Formulation of Bayesian Analysis for Games with Incomplete Information. *International Journal of Game Theory*, 14(1), 1–29.

Milgrom, P. R. (1981). Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics*, 12(2), 380–391.

Milgrom, P. R. & Stokey, N. (1982). Information, Trade, and Common Knowledge. *Journal of Economic Theory*, 26, 17–27.

Mirrlees, J. A. (1971). An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies*, 38(2), 175–208.

Mu, X., Pomatto, L., Strack, P., & Tamuz, O. (2021). From Blackwell Dominance in Large Samples to Rényi Divergences and Back Again. *Econometrica*, 89(1), 475–506.

Newbery, D. M. G. & Stiglitz, J. E. (1982). The Choice of Techniques and the Optimality of Market Equilibrium with Rational Expectations. *Journal of Political Economy*, 90(2), 223–246.

Newbery, D. M. G. & Stiglitz, J. E. (1984). Pareto Inferior Trade. *Review of Economic Studies*, 51(1), 1–12.

Pathak, P., Sönmez, T., Ünver, M., & Bumin, Y. M. (2020). Leaving No Ethical Value Behind: Triage Protocol Design for Pandemic Rationing. Working Paper No. 2695. National Bureau of Economic Research. https://www.nber.org/papers/w26951.

Pęski, M. (2008). Comparison of Information Structures in Zero-Sum Games. *Games and Economic Behavior*, 62(2), 732–735.

Phelps, E. S. (1972). The Statistical Theory of Racism and Sexism. *American Economic Review*, 62(4), 659–661.

Radner, R. & Stiglitz, J. E. (1984). A Nonconcavity in the Value of Information. In M. Boyer & R. Khilstrom (Eds.), *Bayesian Models in Economic Theory*. Amsterdam: Elsevier Science Publications, pp. 33–52. Reprinted in *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*. Oxford: Oxford University Press, 2009, pp. 537–554.

Rahman, D. (2012). But Who Will Monitor the Monitor? *American Economic Review*, 102(6), 2767–2797.

Rey, P. & Stiglitz, J. E. (2023). Moral Hazard and Unemployment in Competitive Equilibrium. In *Selected Works of Joseph E. Stiglitz, Volume IV: Rethinking Macroeconomics*. Oxford: Oxford University Press, Chapter 12.

Rey, P. & Tirole, J. (1986). The Logic of Vertical Restraints. *American Economic Review*, 76(5), 921–939.

Riley, J. G. (1975). Competitive Signalling. *Journal of Economic Theory*, 10(2), 174–186.

Riley, J. G. (1979). Informational Equilibrium. *Econometrica*, 47(2), 331–359.

Riley, J. G. (2001). Silver Signals: Twenty-Five Years of Screening and Signalling. *Journal of Economic Literature*, 39(2), 432–478.

Roth, A. E. (2003). The Origins, History, and Design of the Resident Match. *JAMA*, 289(7), 909–912.

Roth, A. E., Sönmez, T., & Ünver, M. U. (2004). Kidney Exchange. *The Quarterly Journal of Economics*, 119(2), 457–488.

Roth, A. E., Sönmez, T., & Ünver, M. U. (2005a). A Kidney Exchange Clearinghouse in New England. *American Economic Review*, 95(2), 376–380.

Roth, A. E., Sönmez, T., & Ünver, M. U. (2005b). Pairwise Kidney Exchange. *Journal of Economic Theory*, 125(2), 151–188.

Roth, A. E., Sönmez, T., & Ünver, M. U. (2007). Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences. *American Economic Review*, 97(3), 828–851.

Roth, A. E., Sönmez, T., & Ünver, M. U. (2008). Kidney Paired Donation with Compatible Pairs. *American Journal of Transplantation*, 8(2), 463–463.

Roth, A. E., Sönmez, T., Ünver, M. U., Delmonico, F. L., & Saidman, S. L. (2006). Utilizing List Exchange and Nondirected Donation through 'Chain' Paired Kidney Donations. *American Journal of Transplantation*, 6(11), 2694–2705.

Roth, A. E. & Sotomayor, M. (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. New York: Cambridge University Press.

Rothschild, M. & Stiglitz, J. E. (1970). Increasing Risk: I. A Definition. *Journal of Economic Theory*, 2(3), 225–243.

Rothschild, M. & Stiglitz, J. E. (1971). Increasing Risk: II. Its Economic Consequences. *Journal of Economic Theory*, 5(1), 66–84.

Rothschild, M. & Stiglitz, J. E. (1973). Some Further Results on the Measurement of Inequality. *Journal of Economic Theory*, 6(2), 188–204.

Rothschild, M. & Stiglitz, J. E. (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics*, 90(4), 629–649.

Rothschild, M. & Stiglitz, J. E. (1982). A Model of Employment Outcomes Illustrating the Effect of the Structure of Information on the Level and Distribution of Income. *Economic Letters*, 10, 231–236. Reprinted in *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*. Oxford: Oxford University Press, 2009, pp. 122–126.

Rothschild, M. & Stiglitz, J. E. (1997). Competition and Insurance Twenty Years Later. *Geneva Papers on Risk and Insurance Theory*, 22(2), 73–79. Reprinted in *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*. Oxford: Oxford University Press, 2009, pp. 160–167.

Roukny, T., Battiston, S., & Stiglitz, J. E. (2018). Interconnectedness as a Source of Uncertainty in Systemic Risk. *Journal of Financial Stability*, 35, 93–106.

Salop, S. (1977). The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination. *The Review of Economic Studies*, 44(3), 393–406.

Salop, S. & Stiglitz, J. E. (1977). Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion. *The Review of Economic Studies*, 44(3), 493–510.

Salop, S. & Stiglitz, J. E. (1982). The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical Agents. *American Economic Review*, 72(5), 1121–1130.

Sappington, D. & Stiglitz, J. E. (1987a). Information and Regulation. In E. Bailey (Ed.), *Public Regulation*. London: MIT Press, pp. 3–43.

Sappington, D. & Stiglitz, J. E. (1987b). Privatization, Information and Incentives. *Journal of Policy Analysis and Management*, 6(4), 567–582. Reprinted in E. Baily & J. Hower (Eds.), *The Political Economy of Privatization and Deregulation*. Aldershot, UK and Brookfield, VT, USA: Edward Elgar Publishing, 1993.

Shapiro, C. (1982). Consumer Information, Product Quality, and Seller Reputation. *Bell Journal of Economics*, 13, 20–35.

Shapiro, C. & Stiglitz, J. E. (1984). Equilibrium Unemployment as a Worker Discipline Device. *American Economic Review*, 74(3), 433–444.

Shiller, R. J. (1981). Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends? *American Economic Review*, 71, 421–436.

Shiller, R. J. (2000). *Irrational Exuberance*. Princeton, NJ: Princeton University Press.

Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics*, 50(3), 665–690.

Spence, M. (1973). Job Market Signalling. *The Quarterly Journal of Economics*, 87(3), 355–374.

Stigler, G. (1961). The Economics of Information. *Journal of Political Economy*, 69(3), 213–225.

Stiglitz, J. E. (1969). Alternative Theories of Wage Determination and Unemployment in L.D.C.'s. Nairobi: Institute for Development Studies, mimeo. (Parts of which subsequently published as Stiglitz (1974c) and (1982d).)

Stiglitz, J. E. (1972). On the Optimality of the Stock Market Allocation of Investment. *The Quarterly Journal of Economics*, 86(1), 25–60.

Stiglitz, J. E. (1973). Approaches to the Economics of Discrimination. *American Economic Review*, 62(2), 287–295.

Stiglitz, J. E. (1974a). Incentives and Risk Sharing in Sharecropping. *Review of Economic Studies*, 41(2), 219–255.

Stiglitz, J. E. (1974b). Theories of Discrimination and Economic Policy. In G. von Furstenberg et al. (Eds.), *Patterns of Racial Discrimination*. Lexington, MA: D. C. Heath and Company (Lexington Books), pp. 5–26.

Stiglitz, J. E. (1974c). Alternative Theories of Wage Determination and Unemployment in L.D.C.'s: The Labor Turnover Model. *The Quarterly Journal of Economics*, 88(2), 194–227.

Stiglitz, J. E. (1975a). The Theory of "Screening", Education, and the Distribution of Income. *American Economic Review*, 65(3), 283–300.

Stiglitz, J. E. (1975b). Information and Economic Analysis. In J. M. Parkin & A. R. Nobay (Eds.), *Current Economic Problems*. Cambridge: Cambridge University Press, pp. 27–52. Reprinted in *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*. Oxford: Oxford University Press, 2009, pp. 29–52.

Stiglitz, J. E. (1975c). The Efficiency of Market Prices in Long Run Allocations in the Oil Industry. In G. Brannon (Ed.), *Studies in Energy Tax Policy*. Cambridge: Ballinger Publishing, pp. 55–99. (Report written for the Ford Foundation Energy Policy Project, August 1973.)

Stiglitz, J. E. (1977). Monopoly, Non-Linear Pricing and Imperfect Information: The Insurance Market. *Review of Economic Studies*, 44(3), 407–430.

Stiglitz, J. E. (1982a). The Inefficiency of the Stock Market Equilibrium. *Review of Economic Studies*, 49(2), 241–261.

Stiglitz, J. E. (1982b). Information and Capital Markets. In W. F. Sharpe & C. M. Cootner (Eds.), *Financial Markets: Essays in Honor of Paul Cootner*. Englewood Cliffs, NJ: Prentice-Hall.

Stiglitz, J. E. (1982c). Self-Selection and Pareto Efficient Taxation. *Journal of Public Economics*, 17, 213–240.

Stiglitz, J. E. (1982d). Alternative Theories of Wage Determination and Unemployment: The Efficiency Wage Model. In M Gersovitz et al. (Eds.), *The Theory and Experience of Economic Development: Essays in Honor of Sir W Arthur Lewis*. London: Allen & Unwin, pp. 78–106.

Stiglitz, J. E. (1984). Information, Screening and Welfare. In M. Boyer & R. Khilstrom (Eds.), *Bayesian Models in Economic Theory*. Amsterdam: Elsevier Science Publications, pp. 209–239.

Stiglitz, J. E. (1985a). Equilibrium Wage Distributions. *Economic Journal*, 95, 595–618.

Stiglitz, J. E. (1985b). Credit Markets and the Control of Capital. *Journal of Money, Banking, and Credit*, 17(2), 133–152.

Stiglitz, J. E. (1987). The Causes and Consequences of the Dependence of Quality on Price. *Journal of Economic Literature*, 25(1), 1–48.

Stiglitz, J. E. (1989). Imperfect Information in the Product Market. In R. Schmalensee & R. Willing (Eds.), *Handbook of Industrial Organization*, Vol. 1. Amsterdam: Elsevier, pp. 769–847.

Stiglitz, J. E. (1992). Banks versus Markets as Mechanisms for Allocating and Coordinating Investment. In J. A. Roumasset & S. Barr (Eds.), *The Economics of Cooperation: East Asian Development and the Case for Pro-Market Intervention*. Boulder, CO: Westview Press, pp. 15–38. Reprinted in *Selected Works of Joseph E. Stiglitz, Volume II: Information and Economic Analysis: Applications to Capital, Labor, and Product Markets*, Oxford: Oxford University Press, 2013.

Stiglitz, J. E. (2000). The Contributions of the Economics of Information to Twentieth Century Economics. *The Quarterly Journal of Economics*, 115(4), 1441–1478.

Stiglitz, J. E. (2002). Information and the Change in the Paradigm in Economics (Abbreviated version of Nobel lecture). *American Economic Review*, 92(3), 460–501. Reprinted in *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*. Oxford: Oxford University Press, 2009, pp. 53–98.

Stiglitz, J. E. (2009). *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*. Oxford: Oxford University Press.

Stiglitz, J. E. (2010). *Freefall: America, Free Markets, and the Sinking of the World Economy*. New York: W. W. Norton.

Stiglitz, J. E. (2013a). Introduction to Part V. In *The Selected Works of Joseph E. Stiglitz, Volume II: Information and Economic Analysis: Applications to Capital, Labor, and Product Markets*. Oxford: Oxford University Press, pp. 695–704.

Stiglitz, J. E. (2013b). Information and Competition. In *Selected Works of Joseph E. Stiglitz, Volume II: Information and Economic Analysis: Applications to Capital, Labor, and Product Markets*. Oxford: Oxford University Press, pp. 20–35. (Originally the Inaugural Lecture presented at All Souls College, Oxford, June 1978.)

Stiglitz, J. E. (2018). Where Modern Macroeconomics Went Wrong. *Oxford Review of Economic Policy*, 34(1–2), 70–106.

Stiglitz, J. E. (2020). The Revolution in Information Economics: The Past and the Future. In K. Basu, D. Rosenblatt & C. Sepulveda (Eds.), The State of Economics, the State of the World. Cambridge, MA: MIT Press, pp. 101–150.

Stiglitz, J. E. (2023). *The Selected Works of Joseph E. Stiglitz, Volume IV: Rethinking Macroeconomics*. Oxford: Oxford University Press.

Stiglitz, J. E. & Greenwald, B. (2003). *Towards a New Paradigm in Monetary Economics*. Cambridge: Cambridge University Press.

Stiglitz, J. E. & Greenwald, B. (2015). *Creating a Learning Society: A New Approach to Growth, Development, and Social Progress*. New York: Columbia University Press.

Stiglitz, J. E. & Guzman, M. (2021a). Where Modern Macroeconomics Went Wrong. *Oxford Review of Economic Policy*, 34(1–2), 70–106.

Stiglitz, J. E. & Guzman, M. (2021b). Economic Fluctuations and Pseudo-Wealth. *Industrial and Corporate Change*, 30(2), 297–315.

Stiglitz, J. E. & Weiss, A. (1981). Credit Rationing in Markets with Imperfect Information. *American Economic Review*, 71(3), 393–410.

Stiglitz, J. E. & Weiss, A. (1983). Incentive Effects of Terminations: Applications to the Credit and Labor Markets. *American Economic Review*, 73(5), 912–927.

Stiglitz, J. E. & Weiss, A. (1992). Asymmetric Information in Credit Markets and Its Implications for Macroeconomics. *Oxford Economic Papers*, 44(4), 694–724.

Stiglitz, J. E. & Weiss, A. (1994). Sorting Out the Differences Between Screening and Signalling Models. In M. O. L. Bacharach, M. A. H. Dempster, & J. L. Enos (Eds.), *Mathematical Models in Economics*. Oxford: Oxford University Press. Reprinted in *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*. Oxford: Oxford University Press, 2009, pp. 223–231.

Stiglitz, J. E. & Wolfson, M. (1988). Taxation, Information, and Economic Organization. *Journal of the American Taxation Association*, 9(2), 7–18.

Teytelboym, A., Li, S., Kominers, S. D., Akbarpour, M., & Dworczak, P. (2021). Discovering Auctions: Contributions of Paul Milgrom and Robert Wilson. *Scandinavian Journal of Economics*, 123, 709–750.

Thaler, R. H. & Benartzi, S. (2004). Save More Tomorrow™: Using Behavioral Economics to Increase Employee Saving. *Journal of Political Economy*, 112(S1), S164–S187.

Veldkamp, L. L. (2012). *Information Choice in Macroeconomics and Finance*. Princeton, NJ: Princeton University Press.

Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1), 8–37.

Weiss, A. (1980). Job Queues and Layoffs in Labor Markets with Flexible Wages. *Journal of Political Economy*, 88(3), 526–538.

Wilson, R. (1985). Incentive Efficiency of Double Auctions. *Econometrica*, 53(5), 1101–1115.

Wilson, R. (1987). Game-Theoretic Analyses of Trading Processes. In T. F. Bewley (Ed.), *Advances in Economic Theory: Fifth World Congress*. Cambridge: Cambridge University Press, pp. 33–70.