

MIT Open Access Articles

Empirical analyses of selection and welfare in insurance markets: a self-indulgent survey

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Einav, Liran and Finkelstein, Amy. 2023. "Empirical analyses of selection and welfare in insurance markets: a self-indulgent survey."

As Published: <https://doi.org/10.1057/s10713-023-00084-3>

Publisher: Palgrave Macmillan UK

Persistent URL: <https://hdl.handle.net/1721.1/151772>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Empirical analyses of selection and welfare in insurance markets: a self-indulgent survey

This Accepted Manuscript (AM) is a PDF file of the manuscript accepted for publication after peer review, when applicable, but does not reflect post-acceptance improvements, or any corrections. Use of this AM is subject to the publisher's embargo period and AM terms of use. Under no circumstances may this AM be shared or distributed under a Creative Commons or other form of open access license, nor may it be reformatted or enhanced, whether by the Author or third parties. By using this AM (for example, by accessing or downloading) you agree to abide by Springer Nature's terms of use for AM versions of subscription articles: <https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>

The Version of Record (VOR) of this article, as published and maintained by the publisher, is available online at: <https://doi.org/10.1057/s10713-023-00084-3>. The VOR is the version of the article after copy-editing and typesetting, and connected to open research data, open protocols, and open code where available. Any supplementary information can be found on the journal website, connected to the VOR.

For research integrity purposes it is best practice to cite the published Version of Record (VOR), where available (for example, see ICMJE's guidelines on overlapping publications). Where users do not have access to the VOR, any citation must clearly indicate that the reference is to an Accepted Manuscript (AM) version.

(prepared for the *Geneva Risk and Insurance Review*)

Empirical Analyses of Selection and Welfare in Insurance Markets: A Self-Indulgent Survey^{*}

Liran Einav and Amy Finkelstein[^]

Abstract: This review article surveys work that has been done using an empirical framework for analyzing selection in insurance markets developed by Einav, Finkelstein, and Cullen (2010). We briefly review that framework, and then describe a number of empirical applications that researchers have undertaken across an array of settings in both insurance and credit markets. We also discuss some of the useful extensions to the original framework that others have made and applied. The review is intended to be useful for scholars who may want to apply the framework in their own work on insurance, credit, or other selection markets.

1. Introduction

A number of years ago, we drew a graph. We had health insurance data and we were trying to use it to estimate the welfare cost of adverse selection and the welfare consequences of potential policy interventions. We were confused. We thought a graph could help us clarify the empirical objects that we needed to estimate. It did.

We were sure someone had drawn this graph before. The model wasn't new – it was essentially the same as in the seminal paper of Akerlof (1970). All that we had done was create a graphical representation that was helpful for showing us how to empirically operationalize that model. We asked a number of colleagues if they'd seen this representation somewhere we could cite. No one had.¹ So we wrote it up along with our empirical estimates on the welfare costs of adverse selection in a health insurance setting (Einav, Finkelstein, and Cullen 2010).

Over the subsequent years, we've been pretty tickled to see that many other researchers have also found this graphical framework useful for their own empirical work. So when Casey Rothschild, the Editor, asked us to write up a survey of work using the framework we developed, we decided to table our qualms about the embarrassingly self-indulgent nature of the exercise and go for it.

^{*} This article was invited by and prepared for the *Geneva Risk and Insurance Review*. It follows up Einav's Geneva Risk Economics Lecture at the annual meeting of the European Group of Risk and Insurance Economists (EGRIE) in Vienna, September 2022. We are very grateful to Casey Rothschild for encouraging us to write this review, as well as for helpful discussion and constructive comments, to Alex Muermann and an anonymous reviewer for helpful suggestions, and to Miray Omurtak for superb research assistance.

[^] Einav: Stanford University and NBER, leinav@stanford.edu; Finkelstein: MIT and NBER, afink@mit.edu.

¹ Well, almost no one. Amy's husband was sure it was in the textbook of Mas-Colell, Whinston, and Green (1995) but, per usual, she ignored him. Several years later when Whinston became her colleague, he also commented that the algebra behind the graph can be found on page 441. Like we said, we knew the idea wasn't new.

In Section 2 we summarize the basic ideas, also referred to (not by us!) as the “Einav-Finkelstein framework” (EF framework in what follows). One of the attractions of the framework is its portability across contexts. As illustration, Section 3 describes a number of different settings where other researchers have applied the EF framework to test for the presence of adverse selection, estimate its welfare consequences if detected, and consider the welfare consequences of various (potential or observed) policies, such as mandates, subsidies, and pricing regulations.

Although there has been a substantial body of work applying the EF framework to the original health insurance context for which we developed it, the applications that we describe also encompass a wide range of other settings. These include other insurance markets – such as crop insurance, flood insurance, unemployment insurance, and disability insurance – as well as credit markets. These papers are all settings in which there is substantial public policy intervention; indeed, adverse selection offers a potential economic rationale for many of these policies. Often, as we discuss, public policy also provides a useful source for the out-of-equilibrium pricing variation that is essential for empirically implementing the EF framework.

Section 4 discusses some of the explicit extensions other researchers have made to some of the specific, simplifying assumptions made in our original framework. These have generated additional conceptual insights and also allowed researchers to apply the framework in settings where the original assumptions seemed ill-suited. In particular, we describe extensions to the basic framework to allow for imperfect competition, for choice over more than two contracts, and for choice frictions, among other extensions.

Many of the papers we describe in Sections 3 and 4 explore the impact of potential policy interventions designed to ameliorate the welfare costs of adverse selection. In Section 5 we illustrate the use of the EF framework for applied policy analysis by discussing several papers that have used the framework and its extensions to shed light on the optimal use of two specific policy instruments that are frequently deployed to regulate private health insurance markets in the United States: subsidies to consumers and risk-adjustment transfers to insurers.

We have aimed our discussion at scholars in the field of risk and insurance who want to (or wonder whether they should) take advantage of the EF framework for their own data and setting. As we discuss the various papers, we therefore emphasize the types of research designs that have been leveraged as well as conceptual insights that may be applicable more broadly, and only briefly touch on the specific empirical findings of each paper.

2. A brief overview of the Einav-Finkelstein (EF) framework

A graphical framework. The EF framework provides a simple and empirically-tractable approach for estimating and analyzing the impact of selection in insurance markets. The framework takes the insurance coverage contracts as given. It is therefore directly in the spirit of Akerlof (1970), as opposed to Rothschild² and Stiglitz (1976) who allow the coverage level to be an endogenous equilibrium outcome.

² No relation to the Editor.

In this section we briefly summarize the approach. A more detailed description for readers not already familiar with the basic idea is available in the original work (Einav, Finkelstein, and Cullen 2010) and in Einav and Finkelstein (2011).

Figure 1 – which we lift from Einav and Finkelstein (2011) – illustrates the key ideas. It shows the demand and cost curves in a market in which a population of potential consumers all face the same binary choice of whether or not to purchase an insurance contract. One can think of that binary choice as one between no insurance and a given insurance product, or between a lower-coverage and a higher-coverage insurance product.

There are two different ways to interpret our assumption that everyone faces the same choice. One is that the insurance provider is not allowed to price discriminate across consumers due to, for example, pricing regulations (which are quite common in the settings we consider). The other is to allow insurance prices to depend on observable characteristics – such as age or gender – and to think of the figure as depicting a narrow segment of the market for which all consumers in that segment are identical on these observables. In both cases, the figure depicts the residual heterogeneity in demand in costs that is unpriced by the insurer. Because firms can only set a single price and cannot distinguish between heterogeneous consumers, optimal pricing is a function of the average consumer who purchases the product, not the marginal consumer.

We assume that the supply side is perfectly competitive, that risk-neutral firms offer a single insurance contract that covers some probabilistic loss, and that risk-averse individuals differ only in their expected insurance claims. We also assume that firm costs are simply the insurance claims that they must pay out; in other words, we abstract from other firm costs, such as marketing their products or processing submitted claims.

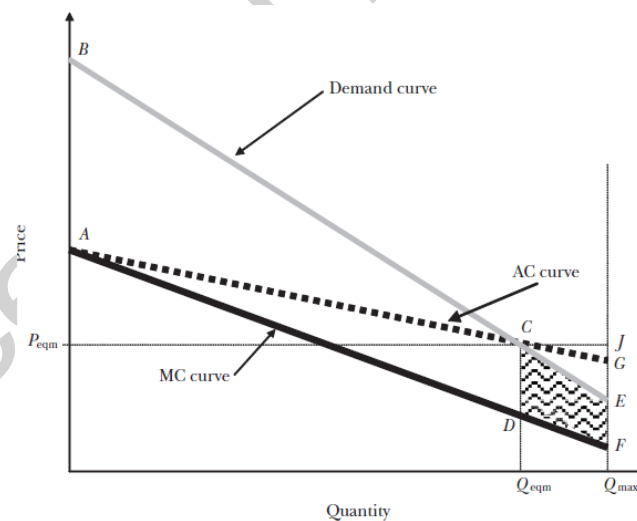


Figure 1: The welfare cost of adverse selection

In this setting, the market “quantity” of insurance (the x-axis in Figure 1) is simply the fraction of individuals who purchase the insurance contract, and firms in this market compete only over what price they charge for the contract. The resulting demand curve in Figure 1 is standard, and reflects (one minus) the cumulative distribution function of individuals’ willingness to pay for the insurance contract.

The distinguishing feature of insurance markets is the cost curve, and, specifically, its link to demand. In many settings, it seems natural to assume that individuals' willingness to pay for insurance is increasing in their expected costs (i.e. expected insurance claims). This is illustrated in Figure 1 by plotting the marginal cost (MC) curve as downward sloping. This represents the well-known adverse selection property of insurance markets: the individuals who have the highest willingness to pay for insurance are those with the highest expected costs; lowering the price brings in marginal consumers who have lower expected costs than those already in the market.

The cost curve in Figure 1 highlights the key distinction of insurance markets (or selection markets more generally) from conventional product markets: marginal cost is driven by demand and thus varies with the selection of consumers. That is, the shape of the cost curve is not driven by the nature of the increasing, decreasing, or constant returns to scale of the production technology. Rather, it is a function of demand-side customer selection.³

The efficient allocation is for everyone whose willingness to pay exceeds their own expected cost to be insured. In other words, the efficient equilibrium is given by the intersection of the demand curve and the marginal cost curve. In Figure 1 this means the efficient equilibrium is for everyone to be insured: given our assumptions, the marginal cost curve is always below the demand curve.⁴ However, because the expected cost of each individual is not priced, the firms must offer a single price for pools of observationally identical (but in fact heterogeneous) consumers. The competitive equilibrium is therefore given by the intersection of demand and average cost (point C in Figure 1), where the average cost curve is defined as the average expected costs of all those individuals who would buy insurance at a given price.

Figure 1 illustrates the fundamental inefficiency created by adverse selection. It arises because the efficient allocation is determined by the relationship between *marginal* cost and demand, while the equilibrium allocation is determined by the relationship between *average* cost and demand. Under adverse selection, marginal costs are always less than average costs. As a result, the equilibrium price is too high and equilibrium quantity is too low, and the welfare loss relative to the first-best outcome is area CDEF, which denotes the lost surplus (i.e. willingness to pay in excess of own expected cost) for individuals whom it is efficient to insure (because their willingness to pay exceeds their own expected cost) but are not insured in equilibrium (because their willingness to pay is less than the average costs of those who are insured in equilibrium). Given the potential for inefficiency, there is a natural potential

³ This stylized example assumes that individuals differ only in their (privately known) expected costs. As shown by Fang and Wu (2018), in a more general setting in which individuals also differ in their utility functions – for example how risk averse they are – it is possible that insurance markets exhibit advantageous selection, in which the marginal cost curve is upward sloping over some portion of the demand curve. This could be the case, for example, when demand for insurance is primarily driven by heterogeneous risk aversion, and the most risk-averse consumers are also the ones with the lowest expected claims. Some empirical examples include the market for private health insurance to supplement the public Medicare insurance in the US (Fang, Keane, and Silverman 2008) and the US market for long-term care insurance (Finkelstein and McGarry 2006). Indeed, insurance brokers quip that long-term care insurance is purchased by the “healthy, wealthy, and anxious.”

⁴ There are a number of ways to relax this, for example by allowing for administrative costs to the firm or moral hazard effects of insurance, either of which could produce a demand curve that crosses the marginal cost curve and an efficient equilibrium in which only some fraction of individuals (those with willingness to pay above their cost) should be insured.

rationale for government intervention in selection markets to try to ensure coverage for individuals whom it is efficient to insure but who are not insured in equilibrium.

From theory to data. One of the key attractions of this framework is that it can be operationalized empirically using existing tools. In particular, as we will see later in this article, for many questions of interest, it is sufficient to know the three curves depicted in Figure 1: the demand curve, the average cost curve, and the marginal cost curve. Estimating these curves is a “familiar” exercise – with familiar challenges that are not unique to insurance markets.

To see this, imagine that the demand and the average cost curves are both linear, as drawn in Figure 1. Then we must estimate

$$(1) \quad D_i = \alpha + \beta p_i + \varepsilon_i$$

$$(2) \quad c_i = \gamma + \delta p_i + u_i$$

where D_i is an indicator variable for whether individual i purchased insurance or not, p_i is the price she faced for that insurance contract, and c_i is the expected cost (to the insurer) of covering individual i .

Estimating demand curves like the one depicted in equation (1) is something that empirical economists have been doing for more than a century (e.g., Working 1927). Estimating demand for insurance is no different. It requires that the economist observe individuals’ insurance options – what contracts are available and at what price – and the choice that they make from their option set. It also requires a research design that shifts insurance prices independently of insurance demand.

Importantly, unlike traditional markets, in which demand curve estimation may take advantage of (exogenous) cost shifters for identification while supply curve estimation would exploit distinct demand shifters, in insurance markets costs are a key component of demand. Cost is determined by the types of individuals who demand the product (at a given price). Therefore, a research design that generates variation in price that is uncorrelated with demand and therefore can be used for estimating the demand curve in equation (1) can also be used to estimate the average cost curve in equation (2).

To estimate the average cost curve in equation (2), the additional data element that is required is c_i , a measure of expected insurance claims for those who purchase the insurance product at various prices. In practice, realized insurance claims are often used to proxy for expected insurance claims, but insurance claims predicted based on unpriced characteristics of the potential buyers can be used as well.

Crucially, the average cost curve is estimated on the *endogenously selected sample* of those who purchase insurance at each price. Exogenous changes in insurance prices affect the composition of people who purchase insurance; this change in composition in turn leads to changes in costs if and only if there is selection. Finally, the marginal cost curve does not require additional estimation. It can be directly computed from the estimates of the demand curve and the average cost curve, since the marginal cost curve is simply the derivative of the total cost curve (i.e. of average costs times demand) with respect to demand.

Knowledge of these three curves is quite powerful. First, from a testing perspective, the slope of the average cost curve provides a clean test of the presence or absence of selection. A downward sloping average cost curve implies downward sloping marginal costs and therefore adverse selection, while an upward sloping average cost curve indicates upward sloping marginal costs and advantageous selection. The inability to reject the null of a flat average cost curve would imply that one cannot reject the null of

no selection, as would be the case in a market with symmetric information. Importantly, the cost curve is estimated on a sample of individuals who have all purchased the same insurance contract. Therefore, this “cost curve test” provides a direct test for selection, in contrast to the earlier and widely-used “positive correlation test” (Chiappori and Salanie 2000; see also Cohen and Siegelman (2010) for a detailed review) which jointly tests for adverse selection or moral hazard.

Second, estimation of the three curves allows the researcher to quantify the welfare loss associated with selection (e.g., the area of CDEF in Figure 1) and the equilibrium and welfare impact of any market intervention (e.g., insurance mandates, subsidies, or risk adjustment). The original theoretical work on adverse selection was motivated by the prospect that adverse selection could impair the efficient operation of markets and open up scope for potentially welfare-improving government intervention. With these three curves in hand, these welfare objects can be empirically estimated.

3. Applications

An appealing feature of the EF framework is that it relies on a standard demand and supply setting, which is a familiar and portable empirical construct. As such, if the researcher has access to the appropriate data and an appealing research design, implementing the EF framework is reasonably straightforward and can be applied across a range of different insurance markets, or any market in which consumer selection plays a major role. This appealing feature stands in contrast to some of our other work on the welfare consequences of selection in insurance markets, such as Einav, Finkelstein, and Schrimpf (2010) in the context of annuities or Einav et al. (2013) in the context of health insurance, which use much more detailed and context-specific econometric modelling that are much less portable across settings.

This is presumably why the EF framework has been used across a range of applications, both within health insurance (which was the context to which we applied it in the original work) and elsewhere. What is required to estimate equations (1) and (2) is that the researcher have information on insurance options and choices, a measure of expected claims for those who purchase insurance, and variation in the price of the insurance options that is orthogonal to demand.

In this section we briefly survey a number of the papers which have made use of the EF framework in its basic form in order to test for adverse selection, estimate its welfare cost, and/or analyze the welfare consequences of public policy interventions. We organize the papers by context, and emphasize in our discussion the key research design elements. Applications that require more substantial extensions of the basic framework are deferred to Section 4.

3.1. An empirical illustration

We begin with a specific empirical application to illustrate more concretely how this framework can be estimated and analyzed. We describe an experiment that Fischer, Frolich, and Landmann (forthcoming) run in which they randomize the price at which they offer hospitalization insurance in rural Pakistan. Prior to this, this insurance was not available to the individuals in the study.

The randomized variation in prices allows them to estimate the demand curve in equation (1). They use baseline survey data to estimate predicted costs for each individual in their sample and then use these predicted costs to estimate the average cost curve in equation (2). The experiment essentially provides the authors with four data points: four (randomly) offered prices, the insurance take-up rate at each of these prices, and the average predicted costs of the individuals who enrolled at each price. The authors then use these four points and impose parametric assumptions in order to estimate demand and average cost curves, and then derive the marginal cost curve. They test for and confirm the presence of adverse selection – a downward-sloping average cost curve – and use the curves to perform welfare analysis.

They conduct this exercise for two different types of insurance contracts, randomizing across villages which insurance contract they offer (in addition to randomizing the price at which it is offered to different people within the village). In some villages they offer insurance coverage contract at the individual level, while in others they offer coverage at the household level, thus covering all members of the household.

Figure 2 illustrates the results from their primary exercise. The left-hand panel shows the estimated demand and cost curves for individual coverage, the right-hand panel shows the estimated curves for household coverage. It shows that the household contract, by forcing members of the same household to pool their risk and jointly decide about their coverage, eliminates most of the adverse selection.

In particular, Figure 2 shows that the marginal cost curve at the individual level (left panel) is much steeper relative to that at the household level (right panel) due to risk pooling. As a result, the implied (counterfactual) share of the market who would obtain insurance in a competitive equilibrium would be more than three times higher if contracts were offered at the household level (rather than the individual level). One particularly nice feature of this exercise is that it illustrates how, in the hands of entrepreneurial researchers who are willing and able to create and offer a new insurance product, the EF framework can be used to assess the welfare cost of selection in a market that does not (yet) exist.




Figure 2: Individual vs. household coverage in rural Pakistan

3.2. Health insurance

Health insurance exchanges. The most common context (to date) for applying the EF framework has been the health insurance exchanges that were operated at the state level, initially in Massachusetts starting in 2006, and then across the United States starting at 2014 as part of the Affordable Care Act (ACA). The regulatory structure of the exchanges – limited flexibility in terms of the coverage offered and strict pricing regulations which offered potential sources of exogenous pricing variation – together with their topical relevance makes them a natural context in which to apply the EF framework.

Hackmann, Kolstad, and Kowalski (2015) use the EF framework to study the 2006 Massachusetts health reform which became known as “Romneycare.” Their paper illustrates that the EF framework can be useful even with relatively minimal data. Specifically, their key research design is the introduction of the health reform, which created a mandate that everyone have health insurance coverage and an associated tax penalty levied on non-poor individuals who did not have health insurance. The authors thus conceptualize the reform as a price change that applied to potential consumers in Massachusetts but not in other states. They use aggregate data on average claims, total enrollment, and insurance premiums in Massachusetts compared to other states, and show how knowledge of these three components before and after the reform allows them to use the EF framework in order to obtain an estimate for the reform’s welfare impact on individuals above 300 percent of the federal poverty line. They estimate that the Massachusetts health reform reduced the average costs of insured individuals, and produced more than \$50 million dollars in annual welfare gains, or just over 4% increase in welfare per person.

Finkelstein, Hendren, and Shepard (2019) also study the Massachusetts health insurance exchanges, focusing on the low-income population that was explicitly excluded from the Hackmann, Kolstad, and Kowalski (2015) analysis. They use administrative data on enrollment and insurance claims provided by the Massachusetts exchange, and a research design that relies on discontinuities in the subsidy schedule that the state offers potential enrollees to generate variation in insurance prices that is orthogonal to insurance demand. They use this setting to estimate the demand and cost curves. They estimate adverse selection in this setting, but their main finding is that for most low-income individuals, willingness to pay for insurance is well below their *own* expected costs. This suggests that adverse selection is not a major factor behind incomplete insurance take-up in this population. Rather, they argue that low willingness to pay stems from the substantial amount of free care that low-income individuals receive if they remain uninsured.

Moving beyond Massachusetts, Kowalski (2014) performs a similar exercise to Hackmann, Kolstad, and Kowalski (2015) in order to study the initial welfare impact of the ACA exchanges in the first half of 2014, and how it varied by states. This exercise illustrates nicely the portability of the EF framework, which allows it to be used as a comparison tool across markets. Kowalski finds that the welfare consequences of the individual health insurance mandate are generally worse in “direct enforcement” states; that is, in states that ceded implementation of the ACA to the federal government.

Using more granular data, Panhans (2019) uses data from the ACA exchange in Colorado during its first two years (2014 and 2015) to test for adverse selection in the newly-formed exchange and quantify its welfare cost. He uses cost data from Colorado’s all-payer claims data, and a research design based on border discontinuities in premiums. Specifically, Colorado was split (by regulation) into 11 rating areas (defined as a collection of adjacent zip codes), and participating insurers were required to offer the

same plans at the same price (up to pre-specified premium adjustments by age and smoking status) within each rating area. The research design relies on comparing individuals who live in adjacent zip codes that are on either side of a rating area border, assuming the individuals are similar (because they live in adjacent zip codes) but face different prices because prices reflect the differences in the entire rating area (of which the particular zip code is negligible in terms of size and importance). He complements this border design analysis with supplementary variation that is driven by the consolidation of some rating areas between 2014 and 2015 (from 11 to 9), and a difference-in-differences variation that relies on differential changes in premiums across areas over the same two years. The primary exercise reveals adverse selection in the form of a downward sloping average cost curve – higher premiums are associated with enrollees with higher average health-care costs and more chronic conditions. The author then draws on estimates from the literature for the demand curve in a similar population and combines this with his cost curve estimates to conduct welfare analysis, finding a welfare cost of adverse selection in this setting of approximately \$250 per person per year.

Further west, Diamond et al. (2021) use data from a financial intermediary from 2012–2015 to study the early years of the ACA exchange in California. Their data contain information about households' monthly enrollment as well as about their medical expenditures. Motivated by the observation that many ACA enrollees drop coverage after a small number of months, the authors use the EF framework to analyze optimal regulation. In particular, they allow a fraction of the households the ability to “strategically” concentrate all their “required” annual medical spending in only several months, and to subsequently drop health-care coverage (and avoid the additional premiums). Using the EF framework, they show that if all adverse selection is driven by such dropouts, welfare would be unambiguously higher if households were required to remain insured for the entire year. However, if adverse selection is driven by “traditional” heterogeneity in risk as well, the welfare implication of regulating dropouts are less obvious. For example, if the overall health-care utilization of enrollees who would choose to drop coverage mid-year is lower than the utilization of full-year enrollees, imposing a minimum enrollment length could cause such individuals to not enroll and thereby increase selection in the market. The authors estimate that in their context requiring a 3.5-month minimum enrollment period maximizes welfare.

Other health insurance markets. Dardanoni and Donni (2016) provide an example of a health insurance application outside of the exchange context. They use the EF framework to quantify the cost of regulation that limits the ability of insurers to offer different prices to observably different individuals. Their primary exercise compares two situations. One in which all consumers in the markets are uniformly priced (due to regulation), so there is only one market clearing price and welfare is computed for the entire market. The second is a situation – which is not allowed in many markets – in which consumers are segmented by observable type, and the insurance offered to each type can be priced separately. In this case, markets are cleared type-by-type, with a distinct type-specific equilibrium price, and a type-specific total welfare, which can then be added up across all market segments (types).

The authors then apply their approach in two markets. One is the market for long-term care insurance, in which the authors use the data used by Finkelstein and McGarry (2006) on insurance options, choices, and realizations of stays in long-term care institutions. The second is the market for supplementary Medicare (aka “Medigap”) insurance in the US, in which the authors use – similar to Fang, Keane, and Silverman (2008) – data from the Health and Retirement Survey (HRSE) on Medigap takeup and

subsequent health-care utilization. In both contexts, the authors find that the pricing regulations lead to substantial welfare loss.

3.3. Other insurance markets

As noted, an attraction of the EF framework is its portability across contexts. To illustrate this, we discuss some of the papers that have applied it outside of the original health insurance context. A number of these papers have studied “markets” that involve publicly provided (and regulated) insurance with traditionally little choice over insurance coverage.

Unemployment insurance. Landais et al. (2021) study unemployment insurance using administrative data from Sweden. In most public insurance systems, employed workers are required to participate in the public unemployment insurance system and are given no choice over the type of coverage that they have. An unusual aspect of the Swedish system is that while there is a mandate that all workers have the basic unemployment insurance package, workers are also given the option to select more comprehensive coverage at a subsidized price.

The authors leverage a government reform that created a sharp increase in the premium for the supplementary unemployment insurance. They apply the cost curve test and find evidence of adverse selection. Yet, as shown in Figure 3, when they apply the EF framework to explore whether mandating the supplemental coverage would be optimal, they conclude that it would not be since there exists a substantial fraction of workers whose willingness to pay for supplementary insurance is below their own expected costs. Instead, they find that the welfare-maximizing policy would be a subsidy that expands the set of people who obtain supplemental coverage, but without leading to everyone having coverage (as a mandate would). This paper illustrates one of the key uses of the EF framework: the ability to go beyond simply testing for the presence (or absence) of adverse selection, and to analyze the welfare consequences of potential policy remedies. In this case, despite detecting adverse selection, the authors conclude that a canonical policy response to adverse selection – a mandate – would not be welfare improving.




Figure 3: Selection in the market for supplemental unemployment insurance

Workers' compensation insurance. Workers' compensation insurance in the US protects employers from tort liability resulting from on-the-job accidents, and provides employees with insurance coverage for medical costs and lost earnings resulting from workplace injury or illness. In all but one state, employers are mandated to provide workers' compensation. In Texas, however, employer provision of this insurance is voluntary (yet heavily regulated). Cabral, Cui, and Dworsky (2022) use administrative data from Texas – and a set of regulatory updates as a source of identifying price variation – in order to apply the EF framework and explore potential economic rationales for mandatory workers' compensation. They find little evidence for adverse selection in this setting, and rule out several other potential economic rationales for a mandate, such as market power and externalities. They conclude that the mandate (which is near universal in the US) is suboptimal, or that it is optimal due to some other reason that they cannot explore, such as choice frictions.

Disability insurance. Cabral and Cullen (2019) apply the EF framework to private long-term disability insurance coverage in the US. They use data from a large employer on its employees' insurance options, choices and claims, and a research design that leverages a sharp change in premiums implemented by the employer during their study period. They find evidence of a demand response to price but lack sufficient power (since long-term disability claims are a relatively rare event) to test for selection. Nonetheless, they can use their claims data to compare their demand estimates to observed claims.

An interesting aspect of their context is that the private disability insurance market that they study provides a supplement to the free, publicly-provided disability coverage available to everyone through the Social Security Disability Insurance (SSDI). The authors therefore use the EF framework to extend their estimates beyond the market from which they are obtained, and derive estimates of cost and willingness to pay for the (free) public coverage, for which direct demand estimation is not feasible (because it is publicly provided to everyone for free). Based on these estimates, they conclude that there would be welfare gains from increasing the generosity of the existing public disability insurance program. In this sense, the paper is similar to the work of Fischer, Frolich, and Landmann (forthcoming), which we discussed earlier, in that the researchers have found a creative way to apply the EF framework to a market that does not exist.

Seibold, Seitz, and Siegloch (2022) perform a similar exercise in the context of the disability insurance market in Germany, which is also a combination of complementary public and private coverage options. They exploit a reform in the public disability-insurance system that reduced the scope of coverage for younger workers, and study the impact of this reform on demand and claims for private disability insurance. Their analysis reveals little evidence of adverse selection in the private market, and willingness to pay that is low relative to expected cost. They present evidence suggesting that the low willingness to pay may be driven by individuals under-estimating their disability risk.

Defined benefit pension plans. Defined benefit pension plans – which pay retirees a fixed monthly amount until they die – provide a form of longevity insurance typically referred to as an annuity. Pensioners with higher expected longevity are therefore more costly to insure. Fitzpatrick (2015) uses the EF framework in the “market” for defined-benefit pension plans for public-school teachers in Illinois. All teachers were provided basic retirement benefits, but had the option to upgrade their retirement benefits by deferring current compensation. The amount required to be deferred for a given upgrade in retirement benefits varied both across the school districts in which the teacher worked and within districts as a function of the teacher's years of experience. Using this variation, the author estimates demand for this upgrade, as well as the projected cost of these incremental benefits. She detects

adverse selection in the market using the cost curve test, but her main finding is that teachers' willingness to pay for this upgrade is much lower than the present value of the benefits it provides, which raises questions as to the way such benefits are structured and marketed to teachers.

Flood insurance. A number of papers have applied the EF framework to the market for flood insurance run by the National Flood Insurance Program. This government program is the source of virtually all flood insurance policies for properties in the US. Interest in this market is prompted in part by the increased frequency of extreme weather events and the fact that, in high-risk flood zones, many properties are not insured. These papers all use administrative data on flood insurance policies and claims.

Collier et al. (2021) and Wagner (2022) both leverage premium variation originating from congressional reforms in 2012 and 2014 that substantially increased insurance premiums for houses built before the implementation of construction standards. Both papers find evidence of a strong demand response to the higher insurance premiums, but little evidence of selection. For example, Collier et al. find that while the average price of insurance increased by 91 percent over a 6-year period, the average claims of those who had insurance increased by less than 4 percent. In addition, Wagner (2022) uses the EF framework to document that adverse selection is unlikely to be a key driver of low levels of flood insurance take-up in high-risk areas as many property owners' estimated willingness to pay is below their *own* expected costs. While the nature of this finding is similar to Finkelstein, Hendren, and Shepard (2019), which we discussed earlier, Wagner present evidence suggesting that this low willingness to pay for flood insurance is likely driven by misperceived information of risk rather than by the expectation of homeowners to be bailed out in the event of a flood.

Mulder (2019), however, finds evidence for a different form of adverse selection in this same market. This selection occurs not on the margin of insurance purchasing as in the standard models, but rather on the margin of selective attrition out of the market. That is, he finds that policyholders with lower expected flood insurance claims are more likely to let their coverage lapse, while policyholders with higher expected claims are more likely to renew their policy. The average risk (and cost) of a given cohort of policyholders thus grows over time due to selection of who renews their policy. Average costs are therefore increasing, not in the price of insurance (as in the standard model of adverse selection in insurance examined by Collier et al. (2021) and Wagner (2022)), but rather in time since the policy has been issued. The author argues that this "dynamic adverse selection" is driven by policyholder learning of their own risk over time. He combines his cost curve estimates with estimates of demand – using premium rate increases that applied to some types of policies but not others – and illustrates the welfare implications using the EF framework.

Crop insurance. Citino, Palma, and Paradisi (2021) analyze the market for crop insurance in Italy. Like flood insurance, interest in this market stems in part from the combination of increasing severity of damages due to climate change and low rates of insurance coverage. In a similar spirit to the Collier et al (2021) and Wagner (2022) papers, Citino, Palma, and Paradisi use administrative data on insurance purchases and insurance claims and a research design that relies on an EU-wide reform which lowered the cap on crop insurance subsidies, thus raising prices sharply. They find that the increase in premiums reduces demand and increases average claims, consistent with adverse selection. In a separate analysis, they also show that insurance take-up responds quite strongly to extreme weather events, suggesting that choice frictions (due to learning and imperfect information about climate risk) may also play an important role in explaining demand in this market.

3.4. Credit markets

While the EF framework was originally developed in the context of an insurance market, it is also applicable to any other “selection market.” That is, to markets in which the sellers’ costs depend not only on how many units of their products they sell, but also on which customers are buying it. Einav, Finkelstein, and Mahoney (2021) provide a more comprehensive discussion of the distinguishing features of selection markets, and how they may impact empirical analyses.

Arguably, the most common selection markets outside of insurance are in credit markets. Here, the price variable is the interest rate rather than the insurance premium, and the cost variable is the expected loan repayment rather than the expected insurance claims. Yet, while conceptually similar, adapting the EF framework to the market for loans requires some modifications for welfare analysis. In particular, in insurance markets it seems natural to assume that the price paid for insurance is sunk, and therefore that the only reason for expected costs to vary with price is from how the price affects selection into the market. In contrast, this assumption is less tractable for credit markets. Here, the interest rate is the typical market-clearing price variable, and different interest rates would not only affect selection into who takes out a loan, but also affect the downstream probability of default by a fixed set of borrowers. Therefore, as emphasized by DeFusco, Tang, and Yannelis (2022), while the demand curve for credit reveals the maximum quoted interest rate that borrowers are willing to accept, this demand curve must be “adjusted down” for borrower’s expected default rate (to capture the impact of interest rate on default) in order to be used as an estimate of willingness to pay, which can then feed into welfare analysis.

Xiang (2020) and DeFusco, Tang, and Yannelis (2022) both use data from randomized experiments carried out by lenders to apply the EF framework to testing for selection and estimating its welfare consequences in consumer credit markets. In the context of the US credit card market, Xiang (2020) leverages a randomized experiment in which the credit card issuer randomized the credit card offers across potential borrowers. DeFusco, Tang, and Yannelis (2022) study a randomized experiment carried out by a Chinese internet lending platform in which interest rate offers were randomized across potential borrowers. Both studies use data on the take up of the loan offers and their subsequent repayment history to estimate the demand and cost curves. Both studies find evidence of adverse selection.

Figure 4 shows the demand for loans and the corresponding AC and MC curves estimated by DeFusco, Tang, and Yannelis (2022). As discussed above, because the key market-clearing variable is the interest rate, the cost curves are adjusted (“scaled” in the language of DeFusco, Tang, and Yannelis 2022) to reflect the fact that the quoted interest rate (that is, the price) is not what the borrowers will subsequently pay given the possibility of default. Using their estimates, they find that adverse selection leads to inefficiently high interest rates in the market, of approximately 30% relative to 9% in a symmetric-information benchmark. Yet, this large effect on interest rates translates to only modest impact on welfare because they estimate demand to be relatively inelastic. Xiang (2020) also finds non-trivial welfare costs of selection in his (credit cards) context.




Figure 4: Selection in credit markets

Lieberman et al. (2018) and Jansen et al. (2022) both use the EF framework to study the impact of information deletion on welfare in credit markets. Limits to the information that lenders can use are a common policy tool in consumer credit markets. Lieberman et al. (2018) study a policy in Chile in which information about individuals' past defaults was no longer available to banks when considering making loans to these individuals. Jansen et al. (2022) use data on auto lending in the US to study the impact of the US legal requirement that information on consumer bankruptcies be removed from credit reports after seven years. In both cases, information deletion generates identifying variation in the interest rates at which loans are offered to different (observable) types of consumers, and the authors use this variation – together with administrative data on loans and subsequent repayment behavior – to estimate demand and cost curves. Both studies find that the information deletion reduced overall borrowing. However, Lieberman et al. (2018) find evidence of adverse selection and estimate that the policy reduced overall welfare, while Jansen et al. (2022) find no evidence of adverse selection and an aggregate welfare loss from the information deletion policy that is quite small relative to the magnitude of the welfare transfer it generates from borrowers who had no bankruptcy to borrowers who did.

4. Extensions

The simplicity and “bare bones” nature of the EF framework help make it portable across contexts and easily used and applied. Of course, even a relatively portable framework can require adaptation to the specific context. We have already seen some examples in the previous section. In this section, we discuss some of the work in which subsequent researchers have explicitly explored the consequences of relaxing one of the key simplifying assumptions in our original framework. These have generated interesting and useful insights in their own right, as well as expanded the tool kit available to researchers interested in applying the EF framework.

Imperfect competition. A natural extension of the basic EF framework is to allow for imperfect competition. We made the perfect-competition assumption in the original EF framework primarily for

conceptual reasons. Under perfect competition, if there is symmetric information the market will achieve the first best. As a result, the symmetric-information-with-perfect-competition case provides a natural benchmark from which to examine the welfare cost of selection; any welfare difference between a world with and without private information can be attributed to selection. Yet, this conceptual clarity may come at a substantial cost of realism. The assumption of perfect competition may be a useful benchmark, but it does not seem ideal for many (most?) insurance markets, which tend to be quite concentrated.

In principle, as mentioned in Einav, Finkelstein, and Cullen (2010; page 884), it is straightforward to use the EF under any other specified model of insurance market competition, including models of imperfect competition. Mahoney and Weyl (2017) explored this, and their analysis revealed that market power can have important – and subtle – implications for the welfare consequences of selection. As they discuss, both adverse selection and market power lead to prices that are “too high” – and quantities that are “too low” – relative to the first-best, efficient benchmark. Relative to a symmetric information setting, adverse selection makes the marginal consumer more attractive to cover than the infra-marginal one, thus reducing the incentives of a profit-maximizing firm to raise prices. As a result, in the presence of imperfect competition, adverse selection may actually improve welfare by reducing markups relative to the symmetric information case. This is a very nice example of the theory of the second best (Lipsey and Lancaster 1956).

Mahoney and Weyl also discuss a related logic which applies to the case of advantageous selection, in which – as mentioned in Section 2 – the cost curves are upward rather than downward sloping, so that the marginal consumer is more expensive to cover than the average one. With advantageous selection, perfect competition leads to prices that are too low and to market equilibrium quantity that is too high. In this setting, introducing imperfect competition would offset this effect; it would lead to higher insurance premiums, which in this case would increase social welfare (given the “too low” prices that prevail in equilibrium with perfect competition).

Cabral, Geruso, and Mahoney (2018) empirically study selection and market power in the context of the Medicare Advantage (MA) market in the US. This is a regulated exchange in which private insurers compete to cover the Medicare-eligible individuals with private Medicare plans (instead of the traditional, public Medicare insurance). The authors’ primary research design relies on a payment reform that differentially affected the per-enrollee amount that the government pays private insurers for covering Medicare-eligible individuals across different counties. They document an incomplete pass through of this payment reform into premiums, and their key exercise explores whether this incomplete pass-through is a result of selection or market power. The research design allows the authors to follow the EF framework and estimate demand for Medicare Advantage, as well as test for the nature of selection using data on cost. They find limited advantageous selection into MA. They then follow the extension of Mahoney and Weyl (2017) discussed above, and find that market power, rather than selection, is the key driver behind the incomplete pass-through they document.

More than two contracts. The original EF framework was presented in the context of two contracts. These two contracts can be viewed as a choice between no insurance and a (uniform) insurance contract, or as a choice between a more vs. less generous coverage where insurers break even on their incremental costs (that is, the difference in their costs between covering the same pool of consumers with more vs. less generous coverage). In principle, as mentioned in Einav, Finkelstein, and Cullen (2010;

page 894), it is easy to extend the basic ideas to more than two insurance contracts, although doing so would naturally not preserve the graphical tractability.

To illustrate the implications of such an extension, consider an extension to three contracts, which include no coverage, low coverage, and higher coverage. The “standard” EF approach could in principle be extended in a straightforward way: markets would clear through two prices (for low and for high coverage) rather than one, and demand and cost would now have to be characterized as systems rather than curves, as they would depend on both prices.

Once again, researchers have explored the implications of such extensions and found that it yields interesting new conceptual insights. Geruso et al. (2021) consider precisely this extension to three contracts – no coverage, low coverage, and high coverage – and show how considering a third option gives rise to two margins of selection: an extensive margin selection between no coverage and some coverage, and an intensive margin selection across different levels of coverage. They then show how – because the two margins are inter-connected – selection on one margin is linked to selection on the other margin, and how different public policies that are aimed at one margin of selection may have (possibly unintended) consequences for the other margin as well. For example, a mandate that everyone have insurance can “solve” the adverse selection problem on the “no coverage” margin, but – by pushing lower-risk individuals into the market, and having them choose the lower coverage option instead of no coverage – can substantially increase the extent of selection between the low- and high-coverage contracts, exacerbating adverse selection on the intensive margin, and leading to a lower share of the population who end up with high coverage. Thus the net impact of the mandate on the amount of insurance coverage is more nuanced.

Handel, Hendel, and Whinston (2015) consider a similar extension. In their setting, all consumers are forced to buy insurance and have a choice of either a low- or high-coverage contract. Yet the situation is not reduced to a two-contract setting because they require markets to clear on both prices (one for low coverage and one for high coverage), rather than only on the market price for incremental coverage (as in the original EF framework). This requirement, combined with the assumption that the “no coverage” option is not available, means that a Nash Equilibrium may not always exist.⁵ In addition, as pointed out by Weyl and Veiga (2017), the Handel, Hendel, and Whinston (2015) market-clearing assumption can make the impact of adverse selection on equilibrium and on welfare much more pronounced than the market-equilibrium assumption in the original EF framework.

Choice frictions. A key assumption behind the welfare analysis in the original EF framework is that of revealed preferences. That is, that the demand curve for insurance also represents individuals’ willingness to pay for insurance. Yet, as many applied examples suggest (including some that were discussed in Section 3, such as Wagner 2022 and Seibold et al. 2022), choice frictions are common in an insurance context. Examples of choice frictions that have been documented in insurance markets include issues of limited cognitive ability (Fang, Keane, and Silverman 2008), biased risk perceptions (Spinnewijn 2015), inertia (Handel 2013), and limited information about offered products (Handel and Kolstad 2015). Any or all of these can drive a wedge between estimated demand and willingness to pay, complicating welfare analysis.

⁵ This situation is not common in their application, and rarely arises. When it does, they propose to use a Riley equilibrium (Riley 1979) instead, which ensures existence in their context.

Spinnewijn (2017) extends the EF framework to study such situations. The key insight is to realize that the demand and cost curves in the EF framework remain unaffected, but in the presence of choice frictions the demand curve no longer represents welfare, and therefore one needs to add a third curve – which he calls the “value curve” – that correspond to the “true” welfare associated with each marginal consumer. Thus, the positive analysis of the EF framework stays the same, but welfare analysis requires integration under this “new” value curve instead of under the demand curve. Handel, Kolstad, and Spinnewijn (2019) illustrate this idea empirically using a setting in which they survey US employees in a large company about their knowledge of the employer-provided health insurance products they are offered. Their survey questions allow them to obtain direct information about the choice frictions associated with the employees’ observed choices. They then use the generalized EF framework to derive implications of these behavioral frictions for the welfare consequences of various public policies.

Boyer et al. (2020) also illustrate this point empirically in the context of Canadian long-term care insurance, where take-up rates are very low. The authors field a survey of individuals’ perceptions of risks and their awareness of long-term care insurance products. They document substantial information frictions, particularly on the awareness dimension. They then run a stated-choice experiment in Canada, in which they randomize the choice of long-term care insurance contract offered to individuals and ask them whether they would buy it; they also ask about the individuals’ risk perception and their understanding of the product offered. They use their data to plot the demand curve, cost curves, and value curve in their setting. They find no evidence of adverse selection but conclude that choice frictions play a first-order role. In particular, choice frictions reduce equilibrium take-up of long-term care insurance, and generate large welfare losses.

Other extensions. Hendren (2021) makes the important observation that the EF framework evaluates welfare using a static welfare metric, which is viewed from the individuals’ perspective at the time they make their coverage choice. Yet, when risk-types vary over time – as is the case in many insurance markets, such as health and life insurance – one might be interested in stepping back and assessing welfare implications from an “ex ante” perspective, before individuals’ risk-type is realized; this is sometimes referred to as behind the Rawlsian veil of ignorance. As with the choice frictions discussed above, the positive analysis of the EF framework remains the same but normative implications can be very different when evaluated from such an “ex ante” perspective. Hendren (2021) calibrates a model to illustrate the quantitative importance of this point.⁶

De Meza, Reito, and Reyniers (2021) use the EF framework to make an interesting point: that adverse selection may not only crowd *out* efficient trade (as in the standard case covered by the EF framework), but also may crowd *in* inefficient trades. This is shown by allowing the demand and cost curves in the EF framework to be non-linear, and considering a marginal cost curve that initially (for high risks) lies above the demand curve, then crosses the demand curve (from above), and then crosses it again (from below). As a result, it is efficient (under perfect information) to only provide insurance coverage to

⁶ Casey Rothschild made a closely related point in his discussion of Einav’s “Geneva Lecture” in the EGRIE annual meeting in Vienna, 2022. If individuals only vary in their realized risk, and welfare is assessed from behind the veil of ignorance, high-risk types are worse off from an “ex ante” perspective – i.e. they have lower expected consumption. They therefore would have greater welfare weights in any social welfare function. This could alter results regarding optimal public policies that attempt to combat adverse selection. For example, it would lead to a higher optimal uniform subsidy relative to the baseline EF framework. See also Fleurbaey (2018) for a more extensive discussion of this topic.

intermediate-risk consumers. In a standard insurance setting, this situation could occur if the high-risk consumers are nearly guaranteed to have a loss, thereby their uncertainty is minimal and their insurance benefit is low.

Solomon (2022) extends the EF framework to consider two different insurance markets for two potentially correlated risks, and the incentive to offer a bundled coverage contract that would cover both risks with a single policy. He documents how the correlation structure between the two risks and whether selection is adverse or advantageous affect the bundling incentives and the associated welfare.

5. Applied policy analysis: subsidies and risk adjustment

In this final section, we further illustrate the potential uses of the EF framework for welfare analysis of public policies by focusing on a particular policy issue that has been salient in US health insurance policy: the regulation of private health insurance markets. Specifically, we consider two distinct regulatory policies, and the potential inter-connectedness between the two: subsidies to consumers and/or a risk adjustment system that govern transfers to insurers. These regulatory tools have generated several papers that attempt to guide policy using the EF framework and its extension to imperfect competition.

Subsidies are frequently discussed and deployed in health insurance markets as a way of increasing the affordability of the product to consumers. Glazer and McGuire (2017) use the EF framework to analyze optimal subsidies in the context of private Medicare plans – the Medicare Advantage market discussed earlier. Despite the policy appeal of providing uniform subsidies that do not vary across plans – a policy that many of its proponents refer to as a “level playing field” – they argue that in the presence of risk selection, the EF framework implies that a uniform subsidy across plans is unlikely to be optimal because of the (non-uniform) equilibrium effects it would generate.

Risk adjustments are additional payments made by the market regulator (e.g. the government) to insurance plans in return for their covering a (predictably) riskier pool of individuals. These are often motivated by the existence of pricing regulations that limit insurers ability to price discriminate across individuals with observably different costs, as in the ones analyzed by Dardanoni and Donni (2016) which we discussed earlier. Such pricing regulations encourage insurers to avoid covering observably high-cost buyers. In this context, risk adjustment systems are typically seen as a way of reducing such insurer incentives by effectively substituting payments from the market regulator for covering predictably higher-cost enrollees for the insurer’s ability to charge higher premiums to these enrollees.

Several papers have analyzed the impacts of risk adjustment in the EF framework. In the context of managed competition, Layton et al. (2017) point out that different risk-adjusted payment systems are often assessed based on their statistical properties, such as how much of the predictable risk the risk-adjustment payments capture. The authors propose instead to use the EF framework to motivate an assessment of risk-payment systems that instead rely on economic efficiency, and then illustrate this proposed method using simulated data. Glazer, McGuire, and Shi (2018) make a similar argument in the context of risk adjustment in Medicare Advantage, and again illustrate it using the EF framework. Similarly, Layton (2017) uses the EF framework to provide graphical intuition for the way different risk adjustment payments would affect market equilibrium, and then estimates the demand and cost curves

using individual-level data from a large employer. He shows that risk adjustments could provide substantial welfare gains by reducing adverse selection.

In Einav, Finkelstein, and Tebaldi (2019), we apply a related argument, and use the EF framework to study the inter-connectedness between government targeted subsidies and risk adjustment. The equilibrium approach of the EF framework clarifies that these two regulatory instruments – although typically analyzed separately in the literature – are quite related to each other. We show that, under certain theoretical assumptions, targeted subsidies dominate risk adjustment, and we use data from the California ACA exchange to quantify by how much.

6. Summary

We were drawn to the study of selection in insurance markets by the findings of the original theoretical work on selection in insurance markets (Akerlof 1970; Rothschild and Stiglitz 1976). This work highlighted how selection could impair the efficiency of an unregulated market and, relatedly, that there was the potential for public policy to improve welfare relative to the market equilibrium. As empiricists, we wondered how large the welfare costs of selection might be in particular markets, and what would be the welfare consequences of alternative policy interventions. But we struggled with how to empirically estimate the welfare effects described by the theory. Out of these struggles came (eventually) the EF framework.

Of the many applications we've discussed, some have found evidence of adverse selection and others have not. Those that have found selection have typically gone on to use the framework for its primary purpose: to conduct empirical welfare analyses of the consequences of selection and of public policy interventions. The results have underscored the importance of going beyond testing. In some cases, researchers have found that although adverse selection exists, it plays little role in explaining the lack of insurance coverage in the market, or that it does not impose substantial welfare costs. In other cases, researchers have found that even when adverse selection substantially impairs the efficiency of the market, some standard policy tools can produce welfare gains while others cannot. Our hope is that as the literature continues to develop, it not only investigates these important empirical welfare questions in additional settings, but also uses the tools that have been developed to provide new insights for how to optimally design policy in the presence of selection.

References

- Akerlof, George. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84 (3): 488–500.
- Boyer, M. Martin, Philippe De Donder, Claude Fluet, Marie-Louise Leroux, and Pierre-Carl Michaud. 2020. "Long-Term Care Insurance: Information Frictions and Selection." *American Economic Journal: Economic Policy* 12 (3): 134-169.
- Cabral, Marika, Can Cui, and Michael Dworsky. 2022. "The Demand for Insurance and Rationale for a Mandate: Evidence from Workers' Compensation Insurance." *American Economic Review* 112 (5): 1621-68.

- Cabral, Marika, and Mark R. Cullen. 2019. "Estimating the Value of Public Insurance Using Complementary Private Insurance." *American Economic Journal: Economic Policy* 11 (3): 88-129.
- Cabral, Marika, Michael Geruso, and Neale Mahoney. 2018. "Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage." *American Economic Review* 108 (8): 2048-87.
- Chiappori, Pierre-Andre, and Bernard Salanie. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy* 108 (1): 56-78.
- Citino, Luca, Alessandro Palma, and Matteo Paradisi. 2021. "Dance for the rain or pay for insurance? An empirical analysis of the Italian crop insurance market." Mimeo.
- Cohen, Alma, and Peter Siegelman. 2010. "Testing for Adverse Selection in Insurance Markets." *Journal of Risk and Insurance* 77 (1): 39-84.
- Collier, Benjamin L., Tobias Huber, Johannes G. Jaspersen, and Andreas Richter. 2021. "Household Take-up of Subsidized Insurance." Mimeo.
- Dardanoni, Valentino, and Paolo Li Donni. 2016. "The welfare cost of unpriced heterogeneity in insurance markets." *RAND Journal of Economics* 47 (4): 998-1028.
- De Meza, David, Francesco Reito, Diane Reyniers. 2021. "Too much trade: The hidden problem of adverse selection." *Journal of Public Economics* 204: 104551.
- DeFusco, Anthony A., Huan Tang, and Constantine Yannelis. 2022. "Measuring the welfare cost of asymmetric information in consumer credit markets." *Journal of Financial Economics* 146: 821-840.
- Diamond, Rebecca, Michael J. Dickstein, Timothy McQuade, and Petra Persson. 2021. "Insurance without Commitment: Evidence from the ACA Marketplaces." NBER Working Paper No. 24668.
- Einav, Liran, and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Evidence in Pictures." *Journal of Economic Perspectives* 25 (1): 115-138.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets using Variation in Prices." *Quarterly Journal of Economics* 125 (3): 877-921.
- Einav, Liran, Amy Finkelstein, and Neale Mahoney. 2021. "The IO of Selection Markets." In *Handbook of Industrial Organization*, Volume 5(1), edited by Kate Ho, Ali Hortacsu, and Alessandro Lizzeri. Elsevier. Chapter 14: 389-426.
- Einav, Liran, Amy Finkelstein, Steven P. Ryan, Paul Schrimpf, and Mark R. Cullen. 2013. "Selection on Moral Hazard in Health Insurance." *American Economic Review* 103 (1): 178-219.
- Einav, Liran, Amy Finkelstein, and Paul Schrimpf. 2010. "Optimal Mandates and The Welfare Cost of Asymmetric Information: Evidence from The U.K. Annuity Market." *Econometrica* 78 (3): 1031-1092.
- Einav, Liran, Amy Finkelstein, and Pietro Tebaldi. 2019. "Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies." Mimeo.
- Fang, Hanming, Michael Keane, and Daniel Silverman. 2008. "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market." *Journal of Political Economy* 116 (2): 303-350.

- Fang, Hanming, and Zenan Wu. 2018. "Multidimensional Private Information, Market Structure, and Insurance Markets." *RAND Journal of Economics* 49 (3): 751-787.
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review* 109 (4): 1530-67.
- Finkelstein, Amy, and Kathleen McGarry. 2006. "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market." *American Economic Review* 96 (4): 938-958.
- Fischer, Torben, Markus Frolich, and Andreas Landmann. Forthcoming. "Adverse Selection in Low-Income Health Insurance Markets: Evidence from a RCT in Pakistan." *American Economic Journal: Applied Economics*. Forthcoming.
- Fitzpatrick, Maria Donovan. 2015. "How Much Are Public School Teachers Willing to Pay for Their Retirement Benefits?" *American Economic Journal: Economic Policy* 7 (4): 165-188.
- Fleurbaey, Marc. 2018. "Welfare Economics, Risk and Uncertainty." *Canadian Journal of Economics* 51 (1): 5-40.
- Geruso, Michael, Timothy J. Layton, Grace McCormack, and Mark Shepard. 2021. "The Two Margin Problem in Insurance Markets." *Review of Economics and Statistics*. Forthcoming.
- Glazer, Jacob, and Thomas G. McGuire. 2017. "Paying Medicare Advantage plans: To level or tilt the playing field." *Journal of Health Economics* 56: 281-291.
- Glazer, Jacob, Thomas G. McGuire, and Julie Shi. 2018. "Risk Adjustment of Health Plan Payments to Correct Inefficient Plan Choice from Adverse Selection." in Ana Aizcorbe and others (eds), *Measuring and Modeling Health Care Costs*. (Chicago, IL; online edn, Chicago Scholarship Online)
- Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski. 2015. "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *American Economic Review* 105 (3): 1030-66.
- Handel, Benjamin R. 2013. "Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103: 2643-82.
- Handel, Benjamin R., Igal Hendel, and Michael D. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." *Econometrica* 83: 1261-1313.
- Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105 (8): 2449-2500.
- Handel, Benjamin R., Jonathan T. Kolstad, and Johannes Spinnewijn. 2019. "Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets." *Review of Economics and Statistics* 101 (2): 326-40.
- Hendren, Nathaniel. 2021. "Measuring Ex-Ante Welfare in Insurance Markets." *Review of Economic Studies* 88 (3): 1193-1223.
- Jansen, Mark, Fabian Nagel, Constantine Yannelis, and Anthony Lee Zhang. 2022. "Data and Welfare in Credit Markets." Mimeo.

- Kowalski, Amanda E. 2014. "The Early Impact of the Affordable Care Act, State by State." *Brookings Papers on Economic Activity*: 277-355
- Landais, Camille, Arash Nekoei, Peter Nilsson, David Seim and Johannes Spinnewijn. 2021. "Risk-Based Selection in Unemployment Insurance: Evidence and Implications." *American Economic Review* 111 (4): 1315-55.
- Layton, Timothy J. 2017. "Imperfect risk adjustment, risk preferences, and sorting in competitive health insurance markets." *Journal of Health Economics* 56: 259-280.
- Layton, Timothy J., Randall P. Ellis, Thomas G. McGuire, and Richard van Kleef. 2017. "Measuring efficiency of health plan payment systems in managed competition health insurance markets." *Journal of Health Economics* 56: 237-255
- Lieberman, Andres, Chris Neilson, Luis Opazo, and Seth Zimmerman. 2020. "The Equilibrium Effects of Information Deletion: Evidence from Consumer Credit Markets." Mimeo.
- Lipsey, R.G., and Kelvin Lancaster. 1956. "The General Theory of Second Best." *Review of Economic Studies* 24 (1): 11-32.
- Mahoney, Neale, and E. Glen Weyl. 2017. "Imperfect Competition in Selection Markets." *Review of Economics and Statistics* 99 (4): 637-651.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. Oxford University Press.
- Mulder, Philip. 2019. "Dynamic adverse selection in flood insurance." Mimeo.
- Panhans, Matthew. 2019. "Adverse Selection in ACA Exchange Markets: Evidence from Colorado." *American Economic Journal: Applied Economics* 11 (2): 1-36.
- Riley, John G. 1979. "Informational Equilibrium." *Econometrica* 47 (2): 331-359.
- Rothschild, Michael, and Joseph E. Stiglitz. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics* 90: 629-49.
- Seibold, Arthur, Sebastian Seitz, and Sebastian Sieglöcher. 2022. "Privatizing Disability Insurance." Mimeo.
- Solomon, Adam. 2022. "Bundled Risks in Insurance Markets." Mimeo.
- Spinnewijn, Johannes. 2015. "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs." *Journal of European Economic Association* 13 (1): 130-167.
- Spinnewijn, Johannes. 2017. "Heterogeneity, Demand for Insurance, and Adverse Selection." *American Economic Journal: Economic Policy* 9 (1): 308-43.
- Wagner, Katherine. 2022. "Adaptation and Adverse Selection in Markets for Natural Disaster Insurance." *American Economic Journal: Economic Policy* 14 (3): 380-421.
- Weyl, E. Glen, and Andre Veiga. 2017. "Pricing Institutions and the Welfare Cost of Adverse Selection." *American Economic Journal: Microeconomics* 9 (2): 139-48.

Working, Elmer J. 1927. "What Do Statistical 'Demand Curves' Show?" *Quarterly Journal of Economics* 41 (2): 212-235.

Xiang, Jia. 2020. "Essays on Asymmetric Information in Health Care and Consumer Credit Card Lending Markets." PhD dissertation. Pennsylvania State University.

Accepted manuscript