- Job market shpiels by David Miller, d9miller@ucsd.edu

  Here are my interview preparation materials. You should use these as an example of the type of short speeches you can prepare for your interviews. You should not necessarily follow my outlines or my manner of speech.

  - Tips

    When you practice memorizing them, if you end up saying something that sounds more natural than what you wrote, go back and revise what you wrote. Practice with someone who can listen to you and make eye contact. Move your hands while you talk to help you keep a conversational pace and tone -- you don't want to sound like a recording, so notice where it feels natural to pause and make sure you remember to pause there each time.

  - Job market paper
    - 2 minute version

      (summary)

      My job market paper is about repeated games with private information. I allow the players to communicate and to transfer utility using side payments. I study equilibria that are robust when communication can happen in any order, or when the players can spy on each other.

      The kinds of applications I have in mind are situations in which the players can communicate face to face. An example might be a cartel trying to organize its collusion, or a town council deciding how much to spend on public works.

      Formally, what I do is I impose ex post incentive compatibility, or EPIC for short, on every stage game.

      (result 1)

      My first main result is that, under EPIC, efficiency cannot be supported in equilibrium, no matter how patient the players are. It's a kind of anti-folk theorem. And the reason for this, is if the players try to behave efficiently in the first period, then to provide the right incentives they will need to reduce their utility later on. They can do this either by going to an inefficient equilibrium in the future, or by burning money in the present.

      (result 2)

      For another important result, I look specifically at allocation games, such as repeated auctions or repeated trade. For two player allocation games, I prove that the optimal equilibrium never uses an efficient allocation rule. Instead, it is always possible to improve on the efficient allocation by introducing some inefficient pooling.

      I also show how to calculate optimal equilibria using linear programming.

      (application)

      I think there's a wide range of real world situations that these results can apply to.

      As an example, I use them to provide a new explanation for price wars in collusion.

      I interpret price wars as a way to to burn money, which the firms will want to do in order to satisfy their EPIC constraints. My explanation allows for asymmetric equilibria and full observation of prices and quantities. This contrasts with Green and Porter's explanation, which relied on limited observability, which would prevent the firms from using asymmetric equilibria.

    - 5 minute version

      (motivation)

      This line of research that I've been working on is about repeated games with private information. The usual approach is to assume that communication happens simultaneously, and then to look at perfect public equilibrium. In my job market paper, I take a different approach: I look at equilibria that are robust even when communication is not simultaneous.

      (examples)

      The kinds of applications I have in mind are situations in which the players can communicate face to face. An example might be a cartel trying to organize its collusion, or a town council deciding how much to spend on public works.

      (EPIC)

      Formally, what I do is impose ex post incentive compatibility on each stage of the repeated game. Ex post incentive compatibility, or EPIC for short, means that each player must be willing to reveal his information truthfully, even after he has learned all the other players' private information. EPIC has several important advantages over the usual approach. First, it is robust when communication can happen in any order. Second, it is robust when players can spy on each other. Third, it is robust when players do not have a common prior.

      (result 1)

      My first main result is that, under EPIC, efficiency cannot be supported in any equilibrium in the repeated game, no matter how patient the players are. It's a kind of anti-folk theorem. And the reason for this, is if the players try to behave efficiently in the first period, then to provide the right incentives they will need to reduce their collective utility later on. They can do this either by going to an inefficient equilibrium in the future, or by burning money in the present.

Since the players have to do this to provide incentives in every period, the effect does not disappear when they become more patient.

(result 2)

So, this raises the question: If the players can't support efficiency, what can they support? My second main result is really a set of results, about answering this question—about constructing optimal equilibria. I show that an optimal equilibrium is a solution to a static mechanism design problem, and that a solution to this problem exists under very general conditions. I also show that in the types of models typically used to study private information, the optimal solution can be found by linear programming.

(result 3)

For my third set of results, I look specifically at allocation games, such as repeated auctions or repeated trade. For two player allocation games, I prove that the optimal equilibrium never uses an efficient allocation rule. Instead, it is always possible to improve on the efficient allocation by introducing some inefficient pooling. The benefit of this inefficiency is that it reduces the amount of money that must be burned.

Under private valuations, the equilibrium rule is often stationary, which means that the players always use the same allocation rule and never burn any money. This requires a very inefficient allocation rule, because the side payments among the players must balance exactly.

On the other hand, under partial common values, the optimal equilibrium involves some money burning. This is because in order to eliminate money burning the players would have to ignore their private information.

(application)

I think there's a wide range of real world situations that these results can apply to.

I have another paper, which I'm cowriting with Susan Athey, that looks at the problem of repeated trade. But in my job market paper I use my results to provide a new explanation for price wars in collusion.

For collusion with hidden costs, Athey and Bagwell have shown that efficiency is attainable in a perfect public equilibrium if the firms are sufficiently patient. But this efficient equilibrium is not robust, and it will be broken if communication is not simultaneous, or if firms can spy on each other, or if they have different beliefs about the uncertainty that they face. EPIC is particularly attractive in the context of collusion, because the private information is over production costs, which are based on things that can be spied on, like technology and capital assets. And, since explicit collusion is illegal, it is hard for the firms to ask an external institution to coordinate their communication to be simultaneous.

In the context of collusion, I interpret price wars as a way of burning money, which the firms will want to do in order to satisfy their EPIC constraints. My theory allows for asymmetric equilibria and full observation of prices and quantities. This contrasts with the famous theory of Green and Porter, which relied on limited observability, that prevents the firms from using asymmetric equilibria.

So this is the basic idea of my job market paper. I'd be happy to dig into it more deeply if you're interested.

▼Top questions

- What is ex post incentive compatibility? Is it new?

  EPIC is a relatively recent idea in the literature, but it's based on the idea of dominant strategies mechanism, which have been used for many years in work on auctions and public choice mechanisms. Dominant strategies requires you to want to tell the truth no matter what the other players do, whether they lie about their information or not. In the case of private valuations, EPIC is equivalent to dominant strategies. That's because, with private valuations you don't actually care about their true information—it's only what they say that can affect you. But in cases with interdependent valuations you also care about their true information. Since a mechanism can't use their true information--it can use only what they say--it can't provide them dominant strategy incentives except to support fairly trivial outcomes. On the other hand, with EPIC, you get a pretty strong notion of incentive compatibility, but it's still useable for interdependent valuations.

  In addition to this practical justification, EPIC is justified over dominant strategies by two theoretical results. The first is that EPIC corresponds to a mechanism that is interim incentive compatible for every possible probability distribution. The second, which is from a working paper by Bergemann and Morris, is that EPIC corresponds to a mechanism that is interim incentive compatible in the universal type space.

- I'm still unclear about why the players need to burn money.

  Well, first of all I usually use the words "money burning" as a general term to include both free disposal in the current period and reduced continuation utility in future periods, since these are perfect substitutes in terms of providing incentives. To eliminate money burning in this general sense, you would have to make the players' side payments balance out ex post. That is, if we are the two players, then whatever I pay must be exactly what you receive.

  But with EPIC, our ex post payments are pinned down by our incentive constraints: for a particular realization of our types, I might need to pay $10, you might need to receive $5, and that leaves $5 that needs to be burnt. We could add on a function that looks like a lump sum to you, but which depends on my type, and maybe this would give you the extra $5 after this particular realization, but it would not give you the extra $6 you would need if you realized a different realization. Ultimately, if we want to get close to efficiency, we'll always be left with an ex post imbalance at least some of the time. This means we'll always need to burn money in expectation.

- How serious are the price wars you predict? Is your theory testable?

  So it won't be clear how serious the price wars are until someone runs a simulation that's realistic. I did run a simulation that I think is not so realistic -- I just wanted to see what ballpark I was in. I assumed that there were two firms with costs uniformly distributed on zero-to-one, and single unit demand with a reservation price of one. The price wars I got were about 7% in severity, and they occurred every three to four periods. But I'm not really sure what a period means in this context. It really should be the time period over which costs are uncertain, and so we could imagine that prices could actually move during a period. So instead of having a 7% price war for a whole period, you could have a 28% price war for a quarter of a period and it would provide the same incentives. So I hesitate to say that I know very much about how serious the price wars can be.

  As for testability, the occasions on which price wars occur are not easy to characterize. For any given example, I can calculate which cost vectors lead to price wars, but these differ across examples. On the other hand, one strong implication of my theory is that an optimizing cartel will not allocate efficiently, with the object of avoiding price wars. Another implication is that the size and frequency of price wars should not be responsive to the discount rate, which is similar to Green and Porter but different from Rotemberg and Saloner or Abreu, Pearce, and Stacchetti. Some Finally, Green and Porter's theory doesn't make sense in industries where prices and quantities are observable, but my theory does apply. And my theory can be distinguished from Green and Porter if we can observe that the market shares vary asymmetrically on the equilibrium path.

- How did you get the idea for this paper? (A common question!)

  Well, I was working on the repeated trade paper with Susan Athey, and I had been thinking about how to construct an EPIC mechanism for that specific game. I was able to prove that efficiency was attainable, either approximately or exactly, under various assumptions. But I had some trouble figuring out how to approximate efficiency if the monetary transfers had to balance exactly in every period. After working on it, I soon realized that it wasn't possible.

  I was able to draw a graphical proof to convince myself that the result was general, and I was able to write down a formal proof for a restricted setting pretty quickly. That was the easy part; the hard parts were, one, generalizing it to arbitrary signal spaces; two, proving existence; and, three, getting results about optimal equilibria for the allocation games.

- Why does pooling improve on an efficient allocation?

  By "pooling," I mean that whenever both players' valuations fall below some fixed epsilon, the allocation ignores their valuations and simply gives the object to player 1, for example. Since the players can't affect the allocation in this region, they don't need to make any incentive payments. Furthermore, although the pooling does affect the incentive payments the players have to make outside the pooling region, we can exactly offset this by adding on a transfer for each player that doesn't vary with his own report. So the only changes we need to consider are those within the pooling region.

  So, consider a simple case with symmetric private valuations and the uniform distribution. Within the pooling region we give up some efficiency because the wrong player gets the object half the time. So in expectation we lose half the average difference between the players' valuations in this region.

What do we gain? Well, under efficiency the incentive payments took the form of a second price auction, so the player with the higher valuation would pay the lower valuation of the other player. By pooling, we eliminated these payments. In expectation, this turns out to be twice as much as the efficiency we gave up. The geometry of the situation is somewhat complicated, but it turns out that the entire amount of the payments we eliminated goes toward reducing the amount of money that needs to be burned.

With the uniform distribution and linear private valuations, this logic scales up linearly, such that as the region of pooling gets larger you continue to reduce money burning by twice as much as you reduce efficiency. So at the optimum there is actually so much pooling that no money needs to be burned at all. This is also true for nonlinear situations that are not too far distant from this simple case. And, much more generally, under some mild regularity conditions, it is always true at least on a small region near the very lowest types.

- Isn't this equivalent to a one-shot game with a no-subsidy condition?

  Not exactly. One of the main points of the paper is that the recursive mechanism design problem is equivalent to an equilibrium in the infinitely repeated game on the one hand, and equivalent to a static mechanism design problem on the other hand. But this static mechanism design problem ignores the participation constraints. In a one-shot game with a no-subsidy condition, the players would typically want to deviate from the mechanism, such as to avoid paying the required transfer after some realizations. In other words, they require a subsidy in order to participate, and one way to provide this subsidy is to give them the opportunity to play the game again in the future.

Secondary paper

- 2 minute version

  (summary)

  In my paper on foregone invention, I look at a problem that occurs when there is a potential market for a new product that hasn't been invented yet—but there's uncertainty over the level of demand. The basic idea is that there may be a number of different firms who are aware of the possibility of creating this new product, but these various firms will typically not face the same costs for actually developing it. What we end up with is a situation in which nobody enters the market, even though at least one firm could earn a profit by entering.

  (example 1)

  In the paper I build up a general theory of foregone invention, and then I apply it to four different examples. My first example was motivated by Microsoft, and its monopoly over operating systems. Suppose that both Microsoft and some other firm both have an idea for a new software application that runs on Windows. Suppose that Microsoft does not enter at first, because its costs are too high. Then the other firm must decide whether or not to enter. It will look ahead and it'll realize that if demand turns out to be high, then Microsoft will enter ex post by bundling its version of the application together with Windows. This prospect reduces the independent firm's incentives to invent in the first place.

  One of the striking things about this is that it's a problem even for Microsoft. Ex post it wants to enter the complementary market in order to drive up demand for Windows. But ex ante it would want to commit not to enter the market. That way it would encourage the other firm to invent the complement in the first place.

  (example 2)

  I also apply my general theory to three other examples, of which I'll describe just one for now. And this is an example of two symmetric, horizontally differentiated firms who can deter each other from adding new features to their products. This is because a new feature can bring them into closer competition, if it turns out to be in high demand. Here, the driver for ex post entry is the fear of losing market power rather than the opportunity to gain it.

  My results contribute the the literature, which looks at some similar situations but doesn't consider the effect of demand uncertainty on incentives when costs are heterogeneous.

- 5 minute version

  (summary)

  I'm also interested in applied theory, whether it's something as general as collusion or as specific as Microsoft. In my paper on foregone invention, I took an idea based on Microsoft and I generalized it to a wide range of situations involving innovation.

In the paper, I look at a problem that occurs when there is a potential market for a new product that hasn't been invented yet—but there's uncertainty over the level of demand. The basic idea is that there may be a number of different firms who are aware of the possibility to create this new product, but these various firms will typically not face the same costs for actually developing it. What we end up with is a situation in which nobody enters the market, even though at least one firm could earn a profit by doing so.

That is, there could exist some firm with an entry cost low enough that it could profit in expectation by entering the market, but it fails to enter because it fears that if demand turns out to be high, then some other firm will enter the market ex post and destroy all its rents. I call this foregone invention.

(defense)

This effect is subtly different from a typical patent model, in which one firm simply has an idea that nobody else has thought of. In my setup, all the firms have the same idea, they just differ in the cost they have to incur in order to implement it. The uncertainty over demand is necessary because otherwise any firm that would enter after the first firm has already entered would also be willing to enter beforehand.

(motivation/example 1)

In the paper I build up a general theory of foregone invention, and then I apply it to four different examples. My first example was motivated by Microsoft, and its monopoly over operating systems. Suppose that both Microsoft and some other firm both have an idea for a new software application to run on Windows. Suppose that Microsoft does not enter at first, because its development costs are too high. Then the other firm must decide whether or not to enter. It will look ahead and it'll realize that if demand turns out to be high, then Microsoft will enter ex post by bundling its version of the application together with Windows. This reduces the independent firm's incentives to invent in the first place.

Why would Microsoft enter ex post if it wasn't willing to enter ex ante? Under Bertrand price competition, if Microsoft were just any old firm then by entering it would knock the price down to zero and earn zero profits. But of course Microsoft isn't just any old firm; it has a monopoly over Windows. So it can bundle its new application together with Windows, and this makes Windows more valuable.

But one of the striking things about this is that it's also bad for Microsoft. Ex post, after demand is revealed, Microsoft wants to enter complementary markets in order to drive up demand for Windows. But ex ante it would want to commit not to enter. That way, it would encourage the other firm to invent the complement in the first place.

(generalization)

What I did is generalize this idea so that it can be applied to a wide range of situations. My basic model has two firms, and it shows that invention is foregone for an intermediate level of costs. This is because if any firm has very low costs, it will want to enter in the first place regardless of what happens later on. Similarly, if any firm has very high costs, it will not enter ex post even if demand is high, and so it doesn't pose a threat to the other firm. When both firms have intermediate costs, there can be an equilibrium in which neither firm enters even though one or both would like to. One key condition for this equilibrium to exist is that at least one firm must view the entry decisions as strategic substitutes.

(other examples)

So, in addition to the Microsoft example, I look at three other examples. The idea is to illustrate that the basic forces in the model can operate through different channels. For example, in the case of Microsoft it was its complementary monopoly that led to an advantage in the ex post market. In the second example, I show that even if a firm has a monopoly in an unrelated market, it can gain an ex post advantage if it can commit to bundling its products together. Through this commitment, it creates an artificial vertical relationship where one doesn't arise naturally. The firm must be able to commit to bundling prior to setting prices, because otherwise its best response would be to unbundle its products and undercut the market price.

I'll just mention one other example, and this is an example with two symmetric, horizontally differentiated firms, and they can deter each other from adding new features to their products. This is because a new feature can bring them into closer competition, if it turns out to be in high demand. Here, the driver for ex post entry is the fear of losing market power rather than the opportunity to gain it.

(conclusion)

My results contribute the the literature, which has looked at some similar situations but hasn't addressed the effect of demand uncertainty on incentives when there's heterogeneity of entry costs.

Top questions

- Can't this problem be solved by patents?

  Well, sort of. If the first entrant could patent its invention, then there would be no foregone invention. But there's two reasons to caution against patents in this situation, one theoretical and one practical. The theoretical reason is that patents are a pretty blunt instrument, and they can be harmful if they prevent ex post entry that would have occurred in equilibrium. That is, there are cost parameters for which it is an equilibrium for one firm to enter first, and then the other enters if demand turns out to be high. There are also some games in which invention is not foregone for any cost parameters. The ideal system would give a patent only when necessary, but this would require the social planner to know all the private costs of all the firms.

  The practical reason is that in the case of computer software it is often difficult to patent new inventions. One reason is that software applications were not historically patentable, although nowadays it seems like almost anything can be patented. More importantly, James Bessen suggests that the big firms, like Microsoft, Intel, and Cisco, have developed a lot of defensive patents, and if you want to write software you probably have to license something from them. This gives them the leverage to prevent you from patenting or to force you to license your patent to them.

- What policy prescriptions do your results suggest?

  So, in the Microsoft example I look at a number of possible policies, both for a neutral social planner and for Microsoft itself. It turns out that many of the policy prescriptions suggested by my results are negative. The social planner, for example, may think about instituting patents to prevent ex post entry, which solves the foregone invention problem. But patents also prevent efficient entry. Depending on the parameters, a patent policy can be a very bad idea in the model.

  Microsoft may think about buying up firms that invent complements, instead of competing with them. But in cases where invention would have been foregone, the independent firm's outside option is to reject Microsoft's offer, in which case Microsoft will enter and destroy its profits. So the price Microsoft pays will be low, and it doesn't offer much incentive for inventing. In fact, in cases where Microsoft cannot credibly threaten to enter, because its costs are too high, the possibility of merging with Microsoft offers even better incentives, but these are cases where invention would not have been foregone anyway. So although the possibility of ex post mergers improves things generally, it makes foregone invention an even more severe problem.

  One policy that helps in the model with just two firms is to regulate Microsoft's pricing. If Microsoft is not allowed to undercut the independent firm, and is not allowed to bundle its product with Windows, the problem of foregone invention will be alleviated. This system also gives both firms more incentive for incremental quality improvement if Microsoft does enter. Unfortunately, it might induce inefficient entry by third parties, and it also begs the question of how to define the market when products are differentiated.

Additional paper (in progress)

- 3 minute version

  (setup)

  So, in the paper on repeated trade, which I'm coauthoring with Susan Athey, we look at a simple bilateral trading game. One player is a seller, and the other is a buyer. They both have hidden valuations for the object that they want to trade. In a one-shot game, Myerson and Satterthwaite showed that it's impossible to support efficiency even under the weakest notions of incentive compatibility, individual rationality, and budget balance.

  We originally just wanted to show that efficiency is attainable for a range of discount factors, if we repeat this game infinitely. In contrast to the folk theorem, we wanted a result for moderate patience.

  But once we started to work on it, we realized a couple of things pretty quickly. First, we realized that the basic result we had in mind was actually somewhat trivial once you looked at dynamic programming as a mechanism design problem. The second thing we realized is that in order to apply the fairly strong set of assumptions in the Myerson-Satterthwaite theorem, we were implicitly assuming that there was a comprehensive set of institutions to support the trading relationship.

  (results 1 and 2)

  So our main point is actually not about achieving efficiency in the repeated game, but about the relationship between the game and the institutional environment. We look at four cases. So, for our first case, if there is a competitive insurance market, then the players can impose their budget imbalances on the insurer, and if they are patient enough then they can achieve efficiency even in the absence of other institutions. For our second case, if there is a trusted market maker who can structure the transaction by receiving each player's request and then allocating the object, then the players can achieve efficiency if they are patient enough.

(result 3)

Our third result is based on my job market paper, where I showed that when there is a lack of supporting institutions, efficiency is not sustainable. In fact, in a wide range of examples, the best the players can do is to use a stationary posted price mechanism. In this kind of mechanism, the seller puts the object up for sale if its worth to her is less than some cutoff price. The buyer will then buy it if it is worth more to him than the cutoff price. This simple type of mechanism matches well with many of the economic transactions that we observe in reality, although it's not efficient.

(result 4)

Perhaps our most interesting case is that we show if the players have the ability to self insure by storing their money in a joint savings account, then they can approximate efficiency. This case is especially interesting because it is not a true repeated game: the players' account is bounded, so they have to track their account balances as a state variable. Furthermore, we showed that they can approximate efficiency not only from the perspective of the first period, but also that the proportion of periods in which the right player gets the object tends to one at the limit.

The mechanism that achieves this is somewhat complex. It has four equilibrium regimes, and chooses a regime based on the level of the account balance. The complexity stems from the need to account for the incentives when the players are close to switching regimes. When the account is at an intermediate level, the two middle regimes both allocate efficiently, but when the account balance is below a cutoff the players deposit money into the account on average, whereas when it above the cutoff they withdraw money from it on average. This way they tend to move towards the cutoff value. If they get close to either of the bounds, the players switch to an inefficient regime in order to move away from the bound for sure.

(wrap-up)

So, to wrap up, the message that we want to convey is that, one, efficiency can be attained or approximated under several types of institutional environments, and, two, that institutions affect the form of the optimal relationship.

Research interests

- 3 minute version

(current work)

Much of the research I've been working on so far has focused on repeated games with private information. Although much of the work is pretty theoretical, there is an important real-world theme that underlies the way I think about these questions. And that's the role of institutions in supporting economic relationships.

For example, at one end of the spectrum, my job market paper is about the limits of cooperation in an environment without any institutional support. On the other hand, I have another paper that extends the folk theorem to games with private information, under very strong institutional assumptions.

(institutions)

I think that game theory often tries to abstract away from the institutional environment. But there are implicit assumptions about institutions embedded in the way we construct our models and the restrictions we place on equilibrium. Part of the progress I want to make is to call these assumptions out into the open, and consider more carefully what they mean and how we might want to think about changing them.

(repeated trade)

A third paper, which I'm cowriting with Susan Athey, takes some first steps in this direction. We look at a repeated trading relationship, and identify various types of equilibrium restrictions and what types of institutions are implied by them. The examples we consider include a competitive insurance market, self insurance, market makers, escrow, savings accounts, and contract enforcement.

(embedded institutions)

What I'd really like to do, however, is not simply turn these institutions on and off, but look at them embedded in the environment that they support. For instance, I would want to know how can an institution exist and be sustained? Can it be provided privately, or does it need to be selected by the community? What happens if an alternative institution is introduced?

(communal production)

I should mention one more topic that I've worked on a little bit, because it's related to this idea of looking at institutions in context. The question, which was posed to me by Partha Dasgupta, is about ongoing cooperative production, say in a small village in a developing country. Why does it often seem that when globalization reaches this village and begins to offer new opportunities, things often tend to get worse before they get better? So I developed a simple model of communal production in which the villagers bargain over their surplus after every period. This represents the existing non-market institutions that are in place in the village. The result that I conjecture, and which I think is quite striking, is that if the villagers expect new opportunities to arise in the future, and that these opportunities will be open only to some of them—say, to the young, but not to the old—then these expectations about the future will cause communal to production to actually decrease in the present. This result, which I plan to explore further, suggests that the transition process between different sets of institutions is not necessarily going to be smooth.

So that's a summary of some of the kinds of things I'm interested in working on over the next couple of years.