Shane Greenstein:

I'm professor Shane Greenstein and you're listening to the Harvard business school digital initiatives seminar. A premier seminar series that hosts thinkers and scholars who are pushing forward research on the digital transformation of the economy by conducting and connecting with cutting edge leaders, equipping leaders and building community. The digital initiative seeks not just to study, but also to shape digital transformation. To learn more, check out digital.hbs.edu. Okay. So it's a great pleasure to have Matt back as with us a paper with a speech at Alison, Tom Blake and as is our customer. Go around the room and everyone introduce themselves. So, and Matt knows who's here and let's start with that Chris.

Chris:

Christian entrepreneurial management.

Tony:

I'm Tony [inaudible 00:00:54].

Alex:

Alex McKay.

David:

David Yoffie strategy.

Matt Backus:

Mark [inaudible 00:01:01].

Eric Markin:

I'm Eric Markin, work organizational behavior.

Speaker 20:

[inaudible 00:01:08].

Natalia Wright:

Natalia Wright, PhD student and strategy.

Speaker 9:

[inaudible 00:01:15] technology and operations management.

Speaker 10:

[inaudible 00:01:18].

Mike Wills:

Mike Wills, marketing student.

Speaker 12:

[inaudible 00:01:34].

Speaker 13:
[inaudible 00:01:34] marketing.

Shane Greenstein:
In the back.

Anne:
Anne with the digital initiative.

Gram:
Gram, PhD student in business and economics.

Speaker 16:
I'm [inaudible 00:01:34].

Speaker 17:
This is [inaudible 00:01:36].

Mike:
Mike [inaudible 00:01:41].

Williams:
I'm Williams, computer scientist. It took off to alumni of Sloane and Harvard.

Shane Greenstein:
Shane Greenstein from the tech on the operations.

Matt Backus:
Is this me?

Shane Greenstein:
Yeah.

Matt Backus:
Okay. So thank you everybody for being here. This is a fantastic group to get feedback on this one thing I'll say since, I know not everybody is an economist. I would love to get questions as I go. I assure you won't be being rude. I'm used to economists, so just whatever's on my mind, anything really. So this is a paper that's joint with Steve to Dallas. It's comes out of an effort between Steven to Dallas and Nielsen Theresa to create a research labs at eBay.

Matt Backus:

Which ultimately dissolved. But while we were there and we got our hands on some data that let us ask what we thought were some pretty interesting questions. And so, Steve was one of the organizers of that group and then Tom Blake was at the time, a full time economist at eBay. He's now moved to Amazon and you know, Tom was one of the data people there that was a co-collaborator almost instantly on this bargaining work. Do I need to trigger.

Matt Backus:

Okay. I'm going to talk broadly upfront about this research agenda, which is a few papers. And I'll introduce them in a second. But by way of background, I wanted to think a little bit about economics and bargaining and negotiations as a literature. And I wanted to pose this question as bargaining a success story for economics. Economics is a field that used to a lot of success stories and we got sociologists to take econometrics. And we've got political science to start using game theoretic models.

Matt Backus:

And so, as a field, we feel we've been able to get our tools used by other people and I'm not sure that this is true in the negotiations' literature. And I wanted to think about why for a second because there is a big theoretical literature, right? Going all the way back to Nash. There's also a large experimental literature and I think there's two reasons that, I mean to, to measure our success, right? You just have to look at who's teaching negotiations and bargaining and it's typically not economist.

Matt Backus:

Right? And Max and I were just talking about this before the talk and I think there are a lot of reasons, right? One is that at the time when this turf, and this was Max argument, right? That at the time when this turf was being created in the 80s economist hadn't really figured out the behavioral economics was a serious thing yet. And we were still running around with our rational models, which look nothing like the real world, right? In our models, people don't talk. In reality, they do flap their jaws when they negotiate. Right?

Matt Backus:

And so we didn't really have any way of explaining what was happening. I think there's a second story in here for why economics wasn't successful at capturing much of this terrain, which is that... In bargaining, unlike an industrial organization or other areas, we weren't really disciplined by data because there wasn't a lot of data, right? It turns out to be incredibly hard to get data from real world bargaining and negotiation in the field. Right?

Matt Backus:

And so I was very excited then, when I was at eBay and I discovered that a tremendous fraction of what happens on eBay happens on something that looks very much like Rubinstein alternating sequential offers, bargaining. I'm going to tell you a lot about this platform but the idea of our agenda is to try to re-introduce economics in the negotiations and bargaining literature and to do so rather than from a dogmatic perspective or to impose these old dated information models rather to use the data and figure out what we can say in this area.

Matt Backus:

So there's three papers that I'm going to talk about a little bit today. I always get a little bit confused by talks that are segmented, one paper after another. So what I'm going to do is I'm going to talk mostly

about this last paper in the series. But I'll have an aside advertisement for the others as I go. So, broadly the top one is the most complete. This one is accepted and sent out. The second one is still a little bit in progress.

Matt Backus:

And the third one is, as I like to say, being revised for my website. So not technically in any formal sense existing, but it's on the way and that's what we'll talk about mostly today. So that's going to be the narrative for this particular presentation. So when I start this talk, I always love to put up this quote from Vincent Crawford back in 1982, says bargaining broadly construed is pervasive.

Matt Backus:

It's all over the place. He talks about labor negotiations and trade agreements. Today, we might talk about climate change mitigation, right? You only need to consider these examples to realize that the potential welfare gains for improving the efficiency of bargaining are enormous, perhaps even greater than those that would result from a better understanding of micro-economic policy.

Matt Backus:

So with any economics, if this isn't obviously funny too. We always have a little bit of envy that the macro economists get to ask the important questions, right? We like to think of ourselves as micro economists getting questions right. Macro economists are asking important questions. This is one case where at least one person is going to tell me that what I'm doing is more important than macro economics. So I was thrilled.

Matt Backus:

What I want to think about is the relationship between communication and bargaining and already you're going to see him start chipping away this old role of economics where there really was no role in those formal theoretical models for communication and bargaining. Does the question matter? So I think there's two ways to think about this. The one is the broad pitch, which is how should we design bargaining protocols, right?

Matt Backus:

Should we regulate communication when two parties are bargaining? And you could think about shuttle diplomacy and caucusing of negotiations is a way to regulate that communication mediation and arbitration or another way to do that. And so thinking about the role of communication is important from the perspective of this broader, how do we design bargaining question?

Matt Backus:

A little bit more narrowly though and closer to what I'll actually do in the paper is that I think it's immediately relevant for platforms that are going to host bargaining and negotiations. One of my favorite examples is Saatchi art where you can bargain directly with artists, right? Amazon has actually shortly after restarted this research project, Amazon adopted a bargaining feature called make an offer as a limited introduction.

Matt Backus:

Taobao has lots of bargaining on the platform, right? There's an instant messenger service on Taobao so you can bargain extensively. And on eBay, about 10% of transaction volume I'll show you is happening by an alternating sequential offers mechanism. And one thing that's interesting here is that they're all facing this question of how much communication do we allow and how do we regulate it?

Matt Backus:

And Amazon, which has the highest final value fee of these three, right? That's the fraction of the transaction that they take as a platform. Amazon doesn't allow communication at all. Taobao with the lowest final value fee they do everything up front with the listing, they allow free form communication and eBay, which takes the intermediate final value fee between eight and 12%. They allow some limited communication, right?

Matt Backus:

So they see this, I think, the first door to consideration for these platforms has been preventing gray market transactions. People taking the transaction off the site. That's the risk. And so the question is there some benefit, right? Is this enabling transactions to happen that wouldn't otherwise. So I'll talk a little bit about the background on this question. Theory is in economics is pretty conflicted about this.

Matt Backus:

So, tacit bargaining theory here, I'm thinking back to Nash and Rubinstein at Silent Tongue communication. There is no role, right? Now, one way to think about modeling communication is to take whatever game you were using to model bargaining before and just add a Cheap talk stage right before it. Okay? So this would be the Cheap talk modification to bargaining. And if you do that with a Crawford and smell style Cheap talk model, then what you're going to get is that more communication can actually be bad.

Matt Backus:

And this is a result that's due to my colleague at Columbia about her design. And the idea here is that, regulating communication is a form of pre-commitment. And in general we can reach better equilibrium if we introduce some pre-commitment mechanisms. If you take a mechanism design approach instead. So here I'm thinking of the Myerson-Satterthwaite approach where it's just free form or thinking about a bargaining mechanism. You report your value to the mechanism and it determines an outcome.

Williams:

What does bad mean?

Matt Backus:

Bad is in... This a great question. I should be more specific. So how efficient can an equilibrium be? So in general, there's going to be lots of equilibrium, but what does the frontier look like? And it turns out in Valders paper you can actually get better equilibrium through pre-commitment strategies on communication by shutting it down. So this was a paper about weather when you're organizing a firm, do you want to talk to your managers or do you just want to delegate decisions to them?

Matt Backus:

And he was showing the talk can be that in a setting. So, in a mechanism design approach, this is the Myerson-Satterthwaite approach where you think about general mechanisms Pharrellen and, and Joe feral, Mike Gibbons have a nice paper that shows that you can have Cheap talk matter in equilibrium in the sense that it conveys useful information, but it doesn't change the upper bounds for bargaining efficiency.

Matt Backus:

And so it's irrelevant, right? It's like you could construct these equilibrium, but from a design perspective, it doesn't change how efficient the mechanism can be. And then if we start introducing behavioral elements, right, it seems communication could be good, right? So we don't get much guidance from the theory. And so we turned to experiments instead in economics. Outside of a bargaining setting, thinking about communication, how do we rationalize the importance of communication?

Matt Backus:

Well, there's one very nice experimental result which says if you take that Crawford and Sabelle game and that was a game, remember where I'm biased, I'm communicating some information to you so you can make a decision. But I have preferences over what you do, right? And then in that game, there's an upper bound to how informative equilibrium can be, what Kyanweng showed in the laboratories is that people consistently outperformed that upper bound and they are doing it because they communicate too much in some sense.

Matt Backus:

To rationalize with the rational agents and equilibrium. Okay. Now that might be because we're lie averse, right? Or regret averse in some sense. So we can introduce behavioral elements to try to rationalize that. And an alternative framework for thinking about why communication might matter is it might be there's lots of equilibrium and the communication is helping us select one through the pregame.

Matt Backus:

Okay. So there's lots of reasons to think from the experimental literature that it might be useful more specific to bargaining. There have been experiments thinking about the introduction of communication. I promised, I already told Max I put this in every time I give this talk, not just here, but I like to pull these pictures out of a paper. Then Max wrote some time ago that did an experiment in bargaining where they introduced higher and higher bandwidth in communication.

Matt Backus:

So, let me just explain what these pictures are. So this is a case where you have two sided incomplete information so it doesn't know what it's worth to the buyer. Buyer doesn't know what it's worth to the seller and please stop me if I get it wrong. So the 45 degree line is the efficient frontier. You would always want them to trade if the buyer values it more than the seller. Okay? So that's the efficient upper bound. That's pretty straight forward.

Matt Backus:

And they ran this experiment allowing people to bargain first with no communication. The black diamonds here are where they successfully transacted, and the open squares are where they failed. And

there's a couple of things that are remarkable about this. One of the reasons I like this paper so much is because you see this little line here. So this little line just above the 45 degree line corresponds to something called the Samuelson Chatterjee linear solution.

Matt Backus:

So there is an upper bound to how well we can do with incomplete information bargaining, right? And that that upper bound is not equal to the efficient frontier because we have an incentive to lie, right? I have an incentive as a buyer to say it's not worth that much to me. You as a seller are going to say, "Wow, my costs are really high." And that line is going to eliminate some of the marginal transactions.

Matt Backus:

This upper bound is one way to characterize the Myerson-Satterthwaite second best, the most efficient outcome. Okay. And what's so neat about this is how closely their outcome in the silent treatment corresponds to that Chatterjee Samuelson upper bound. We're missing marginal transactions but getting the info marginal ones. Okay. So I thought that was really cool. But then of course the result of the papers that as you introduce communication, it seems you start getting transactions that are outside the upper bound of what a rational model of bargaining right? In the spirit of Myerson-Satterthwaite would predict.

Matt Backus:

So this is what's so cool that it suggests that communication can actually do better than economists using the most general tool, the characterizing equilibrium could predict you could accomplish. So this seems really promising evidence and it also suggests that using the rational toolkit, we really can't explain this or at least using the models of communication that are available to us so far.

Matt Backus:

Okay. There's been lots of followup work. I'm selectively citing a couple, one that points out that orientation and culture, which might not be very apparent in text communication, but more so in face to face communication who play a role. And then there are some experiments suggesting that more adversarial than bargaining. You can actually get negative results on this, although I'm not totally clear on what. I mean, this seems pretty adversarial to me.

Matt Backus:

So, the experimental results are positive. The question for me as an economist then is, if we have vague results in theory are ambivalent results in theory and some positive suggestive evidence and experiments. We'd really like to see this confirmed in the real world. And the problem is of course that it's very hard to get data on bargain. So the contribution of this paper that I'm going to focus on today is getting his data from best offer where we can see the messages.

Matt Backus:

We're going to get a natural experiment from the rollout of this messaging platform to think about the effects of introducing messaging and the likelihood of a transaction. And then there's going to be a last piece and I put C in quotation marks here because it's going to be in a high dimensional mold, machine learning space. But I'm going to argue that you can actually see convergence to an equilibrium for the Cheap talk, equilibrium that we believe is happening in the data.

Matt Backus:

Okay, so, we have set up, let me introduce the platform, which is eBay's best offer. So I like to infuriate my European colleagues. This is, you get a new apartment and you want to go buy a John Michael biscuit. So you go shopping for your John Michael biscuit. And here it is on eBay. You can get these prints or here we have a four and a half thousand. What looks like a doormat, a bone sample. And so this is the search result page on eBay. Okay?

Matt Backus:

What you see under this price is it says, or best offer, okay? That signals that the seller is going to be willing to accept offers below four and a half thousand dollars. Thankfully. So here I've clicked on the view item page, so I've gone to the next step in the buyer flow. I could buy it for four and a half thousand dollars. I could add it to my cart or, and this is the button that corresponds to the availability of best offer.

Matt Backus:

I can make an offer on the platform. Okay? So I click on that and this little pop up appears where I can enter my offer. You'll see importantly there's a little extra button I could click to add a message to the seller. We'll talk more about this soon and if I send that offer to the seller, they're going to have 48 hours to accept my offer, reject it or counter. If they counter, I'm going to get a notification and they will have 48 hours to accept, reject, or counter. And if I counter and so on and so on.

Matt Backus:

And this looks a lot like Rubinstein alternating sequential offers marketing. So you can imagine it was a pretty good day when I realized that this was on the site. So that's best offer on the platform. There are a few minor deviations from Rubinstein alternating sequential offers bargaining. One is that sellers cannot initiate this process. They can only post a listing. Buyers only can initiate.

Matt Backus:

The other is that there is a cap, which eBay is experimenting with extending of six total offers back and forth. They thought of this is a way to discourage buyer harassment of sellers. I'm not really sure that's a salient concern. And so they're thinking about extending them.

Alex:

This isn't available for all products?

Matt Backus:

Yes, it is.

Alex:

It is all products?

Matt Backus:

Yep. I'm not, probably, there might be some like maybe they make exceptions in the market for idiosyncratic, but in general this is available for the whole platform.

Alex:

Uniformly everywhere?

Matt Backus:

Yep.

Alex:

Oh, then it's not seller initiated. Okay.

Matt Backus:

What do you mean?

Alex:

If sellers can say, I don't want a best offer. This is-

Matt Backus:

Yeah, that's right. Yup.

Chris:

And so if I remember the alternating offers framework correctly, I thought the equilibrium was that, all trade happens in the first period?

Matt Backus:

Yup. And so-

Chris:

What are the frictions that you think lead to alternating offers that-

Matt Backus:

We're going to talk about this in just a second because one of those papers is really just giving you descriptive facts. One of those facts is not surprisingly, we don't immediately transact the entire platform, which is something that we can't explain what those models. And obviously the most salient friction is going to be incomplete information and in the Rubenstein setting, you know my value, right? More generally. That's probably not the case on eBay.

Natalia Wright:

Matt. [inaudible 00:19:37] sorry, what about multiple buyers? I think I bought something at eBay once in my life, so I apologize for-

Matt Backus:

No. It's a great question. Totally natural question. So it doesn't happen very often and so I'm basically going to ignore it. The reason I feel okay about that is because if you were a seller and you knew there were going to be lots of buyers arriving, right? Then we know from [inaudible 00:19:58], you should hold

an auction. And generally they do. And so in the data, it's very rare that you see two buyers negotiating with the same seller at the same time.

Natalia Wright:

So in this situation we think we care about here we are not expecting to see more?

Matt Backus:

Yeah. And-

Natalia Wright:

So, [inaudible 00:20:12] had both buy the best yet at this [inaudible 00:20:15].

Matt Backus:

Yeah. And then we should expect them to hold an auction. And empirically that's true. You rarely ever see-

Natalia Wright:

That's the most of what actually happens on eBay.

Matt Backus:

Yeah. What do you mean the bus gap?

Natalia Wright:

No. The auction [inaudible 00:20:26].

Matt Backus:

Oh, good. Thank you for the setup. So, okay, here's what happens most on eBay by year. So you might've actually seen this picture before because this picture minus this blue line is the basis of a paper by John and Lauren. And, I think [inaudible 00:20:43] wrote a paper that said, look, auctions on the platform are steadily declining from 2004 to 2015. Right down all the way to a little bit below 30% this is a transaction volume.

Matt Backus:

Whereas fixed price listings are on the accent and so the hypothesis of their paper is it must be that we are getting less interested or it must be that it's just too much of a hassle to do auctions. Now there are a lot of other things happening in the background here, like a period of biggest change was the introduction of best offer, which changed the composition of what actually gets shown and on the search results page.

Matt Backus:

But this blue line is what we're introducing to the picture and that's the introduction of best offer bargaining. And, one of the things that we think is happening is as people are shifting away from the auction mechanism, they still want some price discovery mechanism. If you're selling that guitar from

your attic and you don't know what it's worth, you maybe don't know what price to set it at if you're not going to hold an auction. Right?

Matt Backus:

Because those aren't getting as much attention in the search results. The other thing you could do is use fixed, sorry, best offer bargaining and start soliciting offers to learn about demand. So that's my hypothesis for why this is steadily increasing during the period when we are extensively switching away from high transactions cost mechanisms, which this is right? Getting offers as a high transactions cost mechanism.

Matt Backus:

So now it's about 10% of eBay sales. This continues to grow and is of serious interest to the company, right? So they're thinking hard about the kinds of questions that we're asking today. Okay. So, I'm going to step... Now that I've introduced the platform, I'm going to do an aside and introduce the other couple of papers that we're not going to talk about in as much depth. One which I apologize to, those of you who have seen this maybe too many times.

Matt Backus:

One is a paper called a On the Empirical Content of Cheap-Talk Signaling: An Application to Online Bargaining. This was the first one we wrote with the data because the first thing that we saw that was very salient when we looked at the data is that if you look at the... So this is the average offer discount. Think about this is the average first offer as a fraction of the listing price. Okay?

Matt Backus:

And on the X axis here is the buy it now price, which is the listing price, the one that you see when you look at the search results page and the view item page. So this was the four and a half thousand dollars, except now it goes up to 600. Okay? So here we just did a scatterplot splitting into unit bins to see what those average first offers look like. And in red, I've highlighted the cases where the asking price was a round number.

Matt Backus:

And you'll see systematically people are getting lower offers, when they use a round asking price. Okay? And so we were the first to observe this. There was a literature that's looked at the use of round numbers in bargaining, right? They've saw this result, they've done it experimentally and concluded, therefore you should never use round numbers in bargaining. Right?

Matt Backus:

And so, being economists and somewhat wedded, not as much as a long time ago, but a little bit wedded to the rational hypothesis. We said, okay, well if that's true, if it's a terrible idea to use round numbers and bargaining, because there's some weird behavioral misperception of round numbers that it must be that nobody uses them because nobody would be so stupid as to give up. It turns out to be about 10% of your revenues as a seller.

Matt Backus:

It's a serious loss. And this is really like eBay, sellers are pretty serious. They go on the forum, they optimize their search. These are experienced sellers doing this as well. So we looked at that and what we see is not that sellers are avoiding round numbers, but they're systematically clumping on them. Right? So now in order to tell that story, we have to also tell a story of why sellers aren't learning how to select proper numbers.

Matt Backus:

Okay? And so I don't know if anybody watched Chris House. It's like once you have to tell two stories as a diagnostician, then maybe that's too much. And so we went looking for a more parsimonious answer and we tried what we thought. And, I think what people who read this paper at first would naturally think is a crazy answer, which is maybe some of that incomplete information theory from the 80s that we teach in first year micro economics to PhD students now. Maybe some of that stuff is right. Maybe there's a signaling equilibrium here, right?

Matt Backus:

And so what will a signaling equilibrium require, what it requires that these sellers know what they're doing when they use a round number price, but they are consciously setting around number price and they know they're going to get lower offers. Now why would they do that? They might do that if there's some compensating differential. Right? And what we looked at was the time to sale and the likelihood of sale.

Matt Backus:

And what we found is that these round number listings were much more likely to sell, right? So now when we can rationalize this with some heterogeneity, right? Maybe that guitar in Shane's attic, Shane's happy to let it sit there. He's not in a hurry. So he uses a precise number and gets a better price and I'm in a hurry, right? So I need to pay my bills. I'm just an assistant professor and so I need the cash now.

Matt Backus:

And so I use a round number price to close the deal sooner, right? And how do we communicate this to buyers? We communicate it through the use of round numbers. Okay? And so that was the point of this paper, not... I don't believe that anybody is consciously thinking this when they set their prices, right? I believe the experimental results that we have heuristics around these round numbers, but those experimental results also suggest that those who heuristics are flexible, right?

Matt Backus:

That we're reacting to our experiences around round numbers and precise numbers, right? And so how do you rationalize that? Well, you might rationalize it by writing down a model of rational expectations, beliefs, and that's what this encodes, right? So I think it's fundamentally consistent with those results around heuristics, around round numbers. But this is... Part of the reason I enjoyed writing this paper is I think it's an example of where you can find complimentary use of the experiments and understanding mechanisms.

Matt Backus:

But also overlaying these models to understand how those mechanisms persist in equilibrium and work on the site. And there's a lot more in the paper and thinking about how you document a signaling equilibrium and a little bit of econometric interest if we designed an ML test around trying to find

arbitrary discontinuities, which then spawned a second paper with a former grad student of mine, SitaPang from Cornell, which is online now.

Speaker 20:

If you [inaudible 00:26:39] regardless of the price of the item. So if you were selling a house [inaudible 00:26:43].

Matt Backus:

Yeah, we actually looked into the housing data using the same data that Steve Levitt and Chet Seaverson used. I think it was Chicago area sales. You see a much smaller effect. But if you normalize it, turns out, here there's tremendous variation around the mean in houses, that's not true. So if you normalize it by the coefficient of variation, the effect is exactly proportional to this.

Matt Backus:

But you do see it in housing. The data is much noisier and it's less convincing. And it didn't make it into the final version of the paper. But we were able to convince ourselves that. I love that.

Tony:

When I bought my house, I said that we should not have a around digit and my real estate agent looked at me like I was an alien.

Matt Backus:

When I moved from Cornell to Columbia, my real estate agent said, "You got a job for this. We all know this" I don't know which way it goes. But she thought she knew this. She said, this is obvious.

Tony:

My agent said," Oh, you're going to be paying a lot more than this. This is just not going to matter."

Matt Backus:

Yeah, that's right. It depends on where you lived. Okay. So that's the first paper that came out on this. And the second paper, which is still a little bit in progress. It's a bunch of bargaining in the field. And here we also have Brad Larson from Stanford as a coauthor. And this is a more descriptive data. Sorry, a more descriptive paper. This is the paper we wrote to try to get more people involved in this literature.

Matt Backus:

And so what we're doing is we're releasing one year of all the best offer transactions we could convince eBay to let us release. It's not selected except on category. And so it's about a hundred million bargaining listings. You can get there through my website or you can go to the MVR, but you can download this without a license Republic use. And we're hoping that we can get more people interested in the field empirical study.

Matt Backus:

I say empirical as an economist. I separate that from experiments. I know that's weird. But, so anyway, please use this data. I'll show you some facts from the data that I think are interesting. Just looking at pictures. We frame the paper is being about thinking about what can we explain with rational models

and what can't we. One of the fun things we were able to do is what we're calling an empirical game tree, let's just draw.

Matt Backus:

That buyer makes an offer accept, counter, decline and if they accept, counter, decline and so on. And how often do we get to these different portions of the game tree? The Rubinstein model, we immediately agree. And yet, 27% of the time people continue to bargain. That's just inconsistent with those complete information models in which there's immediate agreement. So anyway, this shines a little bit of light, but mostly it was just fun to be able to do this.

Matt Backus:

We'd never been able to do this. And the study of bargaining with empirical data before. Here I can show you the histogram of listing prices that show up in the data set. I mean, you probably won't be surprised that it's biased towards cheaper items, but here you see those round numbers showing up where sellers are systematically using them. This are the precise numbers. And then some of the pictures that we generated or just thinking about how things change as you move from cheap stuff to more expensive stuff.

Matt Backus:

The listing price, sorry, the fraction of the listing price that the initial buyer offer comes in at. So listing price on the X axis Y axis is the buyer's offer is a fraction of that seems to be slightly declining, although this is a little bit deceptive. This is not a big descend, this is from 6.15 to 5.54. And the bargain price that we end up at has this strange hump shape. And this is actually pretty surprising because if you look at the sale price instead of the final negotiated price, you don't see that hump shape.

Matt Backus:

And so what's going on in the difference between those two pictures that some people pay four and a half thousand dollars for the basket, right? They don't bother to bargain. Okay? So here is the probability that we bargain at all conditional on bargaining being available on the listing. And you see for really cheap stuff, a lot of people don't bother. And when they do send an offer for really cheap stuff, it's much more likely to be immediately accepted.

Matt Backus:

Why is this? Because bargaining is costly, right? It's taking time and effort to submit these offers and so if it's really cheap you just don't bother. And if it's really cheap and you do make an offer, the seller doesn't bother. I think this is interesting because economists had been writing down the huddles of bargaining typically in a setting where they assume proportional costs of bargaining, right?

Matt Backus:

So this is the Rubinstein ice cream model that some of it is melting away at a proportional rate. And so those costs then don't induce any differences between small-scale bargaining and large-scale bargaining. But in reality we do see differences. Yep.

David:

Is there any sense in which there's a bargaining breakdown here or breakdown here just means both parties walk away?

Matt Backus:

Right. That means both parties walk away. Yeah. That's going to be important for the... Especially once I get back to the communication paper, it's going to make my results conservative in a way that all that'll show you.

David:

Oh, okay. Was that also part of the tree and I didn't quite understand it.

Matt Backus:

Yeah. Oh, I can go back. Cool. Here we disagree. And sometimes, right? So, the seller has declined your offer. Sometimes you're going to come back and make another offer to the seller. About 37% of the time and about 63% of those 40% of the [inaudible 00:31:49] declines. We just walk away. So these are the walkaway States.

David:

Oh. What does that. The aggregate two as a percentage of total-

Matt Backus:

I think it ends up being in total about, I want to say about a half. I forget for this data set. Oh, yeah. There's lots of walks.

David:

Quite a bit.

Matt Backus:

Yep. I'll show you the summary stats for the communication. I forget off the top of my head.

David:

Oh, yeah. I guess one could rationalize either side came in thinking more optimistically than was true. Yeah, sure.

Matt Backus:

And it could be that it's not efficient for us. You might like your guitars more than I can-

David:

Could be there's no price which I can transact. That's perfectly reasonable.

Matt Backus:

Yeah. Okay. So one of the novel things that we can do with this data set, because we don't only observe bargained outcomes, most of the data we have on bargaining is just outcomes, right? Hospital insurer negotiations. We just see outcomes, not the bargaining process. So one of the things we wanted to do in

this paper is to actually dig into the process of bargaining. But that's really high dimensional data, these back and forth of offers.

Matt Backus:

So here's the trick we introduced to try to at least tame the data to be able to give you descriptive facts. What we said is, look, for every round of bargaining, we're going to create this parameter called gamma. Okay? Gamma is going to be take the last two rounds of bargaining. The convex combination of those last two offers. You can always express the next offer as a con because it's going to be somewhere in between.

Matt Backus:

So you can always use gamma to characterize an offer at time T as a combination of times T minus one and T minus two. So let me get a little bit more specific. In the first round we're going to say what was the asking price. Gamma is the fraction of that that the buyer offers in the next round. So gamma two is going to be the gamma weighted combination of the initial asking price and then the buyer's first offer and so on and so forth.

Matt Backus:

So we think about gamma as a... We call it the cave rate. It's how much did you give in to my last offer? If gamma is zero. And you just repeated your old offer. If gamma is one, then you just said yes, I will go immediately to your offer. Okay? All right. So now that we have those gammas we can start to generate pictures of this evolution, generated by bargaining. And there's two things that we did.

Matt Backus:

One is we just wanted to know how do these offers evolve over time? Right? Sorry, I actually didn't do this one in terms of the gamma's shoot. This is just in terms of the offer amounts. The thing that was interesting here, we're calling this [Cosium dynamics 00:34:18], is that... So this is for a bargaining sequence that lasted six turns and ultimately ended in agreement.

Matt Backus:

What you see is that sellers and buyers are getting closer and closer to each other. Now that might seem perfectly obvious they agreed, so they must be getting closer and closer. But if you look at those old economics bargaining models, that's not how it works. In an economics bargaining model, right? It's a war of attrition as in the [Crampton model 00:34:42] or it's a screening model.

Matt Backus:

But basically the second you give me information, you tell me, okay, I'm actually willing to come down to my price. I capitalize on that. Right? So the question is, if you make a concession, do I respond with a concession, which is what we see in the data? Or do I capitalize on it? Because I know that you are now the high value type and I can extract a lot of rents from you.

Matt Backus:

The ladder is what the economic theory predicts. The former is what we see in the data. And so we call this phenomenon and I don't think we're the first to coin this, but reciprocal gradualism, right? That buyers and sellers are slowly converging to each other. And this is a phenomenon that's really hard to

explain in an economic model. So there's one feature that you can't explain with economics the other and now I really am going to use gammas.

Shane Greenstein:

That [inaudible 00:35:29] did they end up at the halfway point between their two opening offers?

Matt Backus:

Yeah. At the end of this. So let's see. So this is the-

Shane Greenstein:

Because that's what the situation says you out get.

Matt Backus:

Oh, you set me great here?

Shane Greenstein:

Yeah.

Matt Backus:

Let's do those gammas. Because this is the halfway story-

Shane Greenstein:

And then the other question that comes is, are all of these always taking place between strangers? And they always arm's-length in the sense that they don't know each other?

Matt Backus:

Yeah. I'm assuming that.

Shane Greenstein:

Yeah. You're assuming that. But if they're repeated interactions in a sense, which is possible on this platform-

Matt Backus:

So I don't remember the fraction of times that it happens. I remember it's rare. I remember I looked at it and decided I wouldn't worry about it-

Shane Greenstein:

Because you interpreted very differently if they were repeated.

Matt Backus:

That's right. I mean, part of the reason we liked this platform so much is that I don't think that that's happening so much unlike, if I were looking at buyer, supplier contracts where, you know, it's repeated, right? And we know that the people that negotiate those contracts get their bonuses based on revenue

instead of profit, which economists upset. But yeah. So what does this, this is a, remember those Gamma's?

Matt Backus:

This is a histogram of the Gamma's for each round of bargaining. So on the top left, this is the buyer's offer is a fraction of the seller's asking price. See most of the mass is off to the right next here is the sellers counter to the buyers offer conditional on getting to that point in the game tree. Right? Here's the buyer's counter to that and the sellers counter to that. And so there's a few things that are salient here.

Matt Backus:

One is that it seems some information is being revealed. In the sense that we are making concessions early on and less as we move on in the bargaining process. The mass is moving to the left, but to the really stark thing to us was all this mass on 50%. Right? Which is really hard to explain in a theoretical model of bargain. Yeah mike.

Mike Wills:

So, zero means and they sent back a response, they just repeated their offer?

Matt Backus:

Yeah.

Mike Wills:

And is that more effective than not responding?

Matt Backus:

That's a great question. So if this is buyer, seller, buyer, seller, so yes. In the sense that as a seller this continues. So I can't as a seller re initiate this bargaining process.

Mike Wills:

I see. So if I'm not moving from my price. If I don't respond, they can't send me another offer?

Matt Backus:

They can, but you can't send them another offer. So, if you're keeping it alive. That's my guess at this. The other thing that could be happening here is you might be sending a message rationalizing the price. Which you couldn't do if you just declined.

Mike Wills:

That's an interesting empirical test.

Matt Backus:

Yeah.

Mike Wills:

[inaudible 00:38:16].

Matt Backus:

Yeah.

Mike Wills:

Lets find out whether-

Matt Backus:

I don't know off the top of my head. But we should look-

Mike Wills:

[inaudible 00:38:21]. It just always made sense to [inaudible 00:38:23] your last offer.

Matt Backus:

Yeah, we haven't done it, but we should, I think that'd be fun. Okay, so that was the stark result there. Is this settling at 50%. Yep.

Mike:

I'm curious about the scatterplot of gamma one versus gamma two in that... Do you sellers just have, here's what my second offer or first bargaining offers going to be. I'm just curious like the dynamics, because it seems there's some symmetry between those two.

Matt Backus:

It's actually funny because right now we have an RA generating exactly that. That's the more Mike... The more detailed version of this plot. Right? You want to know is the variation in this correlate with variation in this?

Mike:

Yeah.

Matt Backus:

That's, actually, we haven't already created [inaudible 00:39:17]. So I will email it to you.

Mike:

Okay.

Matt Backus:

Yeah. Okay. Oops. Shoot, so I left off the one other slide. The one other thing that we find around these split, the difference offers that we thought was interesting is that if you just do a local polynomial fit to the likelihood of success for any of these, right? And add a dummy for being at 50%, it seems these offers are more likely to be accepted. Right? Even more likely than an offer that is more advantageous to the other party. Right? And this is something that we find very hard to reconcile here. Yup.

Tony:

Have you got anything to link this to product characteristics. So I might imagine that products that have a higher likelihood of ex post moral hazard have very different bargaining dynamics.

Matt Backus:

Yeah.

Tony:

So one of the reasons that you might think that wage bargaining or something like that happens relatively rarely or that wages are somewhat rigid is because of the incentive effect later on. And I don't know if you could try to link this to product reviews or transaction based reviews or something like that to give you a sense of, what the change in incentives in terms of the split of the rents with would do the actual economic outcome.

Matt Backus:

Yeah. No, this is a great question. So this is one of the basic limitations of the data set is that, well for any given product, you could find the product listing and look pretty carefully at what it is. eBay doesn't have a very good standardized way to look at product identity. So in the data set, one thing, we do release this for the cases where they do have a standardized product identity where they have a product ID.

Matt Backus:

Then we can link it to other listings and try to figure out what the average sale price was. And we included that in the data set. But in some sense that's getting away from the interesting set of stuff to study, right. If there's a thick market so that you can control for unobserved heterogeneity, then there's no reason to bargain it.

Tony:

Yeah.

Matt Backus:

Right? And so the extent to which we're interested in this question is the extent to which I can't do very well. I'm thinking about product characteristics. So this is I think one of the struggles for this data set and for this. Yup.

Eric Markin:

It was their thing maybe that 30% percent rule is so usually with so many of the chief products, right? If people, so low cost of bargaining it's only $10, I'll just give him five and see if they take it.

Matt Backus:

Yeah.

Eric Markin:

Just a low mental effort to just throw up 30% set?

Matt Backus:

This is a good question. So, and there's another version of this question which comes out of some work by Devin Pope, which is... And I haven't done this yet. This is on our list as well. If gamma at 50% is near a round number, is that an attractor? Are we more likely to use it? Yeah, I like that. That's something we should consider and haven't done. Yup. Okay. That said. That data set is not what we're going to use today.

Matt Backus:

So that data set is the data set I want you to use and download and write papers with. The data set I'm going to use today is from not eBay.com but ebay.de. So it turns out that in different countries it is a functionally different website. Okay? And so this is the best offer pop-up. Remember you did a search and then you clicked on you item and then you clicked on make an offer.

Matt Backus:

This is the pop up on ebay.de and you'll notice there are a couple of things that are different. One is that it's written in gibberish that I can understand it, it's in euros, but two here I've highlighted where before there was a little plus and then add message to seller. It would say, yeah. It doesn't have that. Okay. And so we, we discovered this early on in the bargaining research agenda.

Matt Backus:

And so we went and asked some product managers, why is it that you don't allow communication on eBay.de and all we got in response were a bunch of jokes about Germans. There was no institutional memory. From our perspective, it appears that this was literally just an accident of the implementation of the eBay site or maybe the product manager rolling out eBay.de just had different priors about the effectiveness of this feature.

Matt Backus:

But in any case, there was no communication for eBay.de you could not send a message accompanying the offer. And by the way, in the US these are 250 character text messages that accompany any given offer. So of course-

Shane Greenstein:

[inaudible 00:43:31] we trimmed the rest of Europe, just Germany?

Matt Backus:

Just Germany. Yup.

Shane Greenstein:

Austria?

Matt Backus:

Yup. I forget whether there is a-

Speaker 9:

Even in German has a weird history though. So started by these two brothers. One of them is here in HPS. Basically they literally just copied him in this really, crass form of plagiarism and their business model for many other tech companies was this, they saw eBay taking off in the States. They're like, "We gotta do this for Germany." And then they got eBay to buy the German eBay and made billions of dollars on this. And then went on to... So Rocket internet is that company around starting tech companies that are literal copies to the point of an ambiguous plagiarism of US websites and then selling them often times to the US anyhow, even Germany-

Shane Greenstein:

[inaudible 00:44:22]-

Speaker 9:

Sorry.

Shane Greenstein:

[inaudible 00:44:23]-

Speaker 9:

Their are two brothers. This was long before my time. There are two brothers both German. One of whom did spend time at [inaudible 00:44:33] business school. Even Germany does have unique DNA relative to the rest of the [inaudible 00:44:43] nation. That's like just a fun fact. HBS,

Matt Backus:

I'm going to [inaudible 00:44:51] more about this later, this is [inaudible 00:44:52].

Williams:

The [inaudible 00:44:53] is Samwer s-a-m-w-e-r brothers-

Matt Backus:

I'm not asking about the-

Williams:

We can talk about it later, but anyway, Germany is like weird in a bunch of ways. I think if I it might have something like [inaudible 00:45:07], that are not-

Matt Backus:

Not originally but, yeah-

Anne:

I'm Anne [inaudible 00:45:13] I should ask what extent is eBay in European countries? because I know that for example, their ad placements or the lack of that is definitely the difference between [inaudible 00:45:24].

Matt Backus:

It can be super different, right? Because a lot of their expansion outside of the US was by buying companies. Right? And it might look like more like a classified company in some countries and more like eBay and other countries and different markets are going to be stronger or weaker. So it can be quite different across countries. I know in France it's very different from the US model. In the UK it's relatively similar.

Matt Backus:

But yeah, there's a lot of heterogeneity, which is part of why you'll see we didn't use in our design a comparison across, European countries. So we found this out and so we of course, begged them to run an experiment, right? Because this is half the reason we had this research labs, right? It was to do some experiments in the firm. And, we said this could be potentially really large gains for the platform.

Matt Backus:

And they said, no, we can't let the Germans talk. So we persisted and trying to convince them that this could be really great. We ran some OLS regressions, we put up some tables, there were even standard errors and then one day they decided that this is brilliant and we should immediately implement it. So they went from this a terrible idea to this brilliant idea with no intermediate range of uncertainty.

Matt Backus:

During which they might have run an experiment. Right? And I had a separate paper about this, but yeah, no, they never doubted their conviction and so they just turned it off. So May 23rd, 2016 they just turn on communication in Germany, not doing the category by category experiment that we were really hoping for. What else is happening in Germany at this time? It turns out there's lots of holidays in may.

Matt Backus:

And importantly, there were no other major site changes at the same time. This is important because, when you see you May papers that are studying site changes, you should be suspicious because they tend to roll out in batches. This was enough. Okay? And what works for us is the following. Remember when you visit a website, you're literally just downloading the website, right? But when you use an app, you're using software that's installed on your phone, right?

Matt Backus:

And in order to update the platform, they would have to rewrite that software. So they updated the website, they did not update the mobile app in Germany. And that's going to be the basis for our different depths. So originally we thought about making comparisons across European countries. That turned out to be a terrible idea because they're so different. But this is a comparison that we're going to get some mileage out of.

Matt Backus:

Is mobile versus desktop users before and after. Okay. I'll do this now because I'm not going to do it later. You must be curious. What in the world are these people saying here? I've used Google translate, so I apologize in advance that some of the messages seem particularly awkward. It's less awkward than it would be if I tried to translate them. They say every... And this is true on the American website as well.

Matt Backus:

People using these messages say everything you can think of, things that seem perfectly rational or what an economist would think would be the rational story. I'm an expert in this and I've traded this for 10 years and you're charging a little bit too much and here's some other listings. Two things like, I want to give it to my dying mother. It's what she's always wanted. Can you give me half off? Right?

Matt Backus:

Things that are obviously not consistent with a rational model. And so since the hardest asking, it's all right. Rarely less, rarely laughed so much. A good reveal. So people say a lot of stuff and there's not an obvious coherent way to interpret it just by reading, these messages. Okay. But for the bulk of the paper, until I get to the end, I'm going to ignore the content of these messages and focus on whether one was sent and I'll come back with a little bit of a machine learning approach to think about the content of these messages at the end.

Matt Backus:

That was the fun slide. Check. Can't get away from it. Okay. So here on the X axis we have time. On the Y axis we have the fraction of buyer, seller interactions with a message. Now one quick aside, before I dig into this guy... My unit of observation is going to be an interaction which is a buyer item ID pair. Okay?

Matt Backus:

So it's the whole sequence of offers in that buyer item ID pair. But that's going to be my unit of observation and I'm going to be ultimately interested in predicting whether they're successfully going to transact. Okay. So here is the fraction that include a message. There's two things. One adoption was instantaneous and I think it's because in Europe they might be used to using ebay.com which is a larger marketplace.

Matt Backus:

And so it might've been natural to see this messaging platform. Here, I've split the sample by whether the buyer is a desktop user or the seller is a desktop user. I'm only going to see this difference for buyers, not for sellers.

Shane Greenstein:

And you know the ID?

Matt Backus:

Yup.

Shane Greenstein:

So you know, if a buyer moves from desktop [inaudible 00:50:00] and back?

Matt Backus:

Yeah, we don't see that. So we see that a little bit between transactions. We don't see this very sorry, interactions. We don't see this much within-

Shane Greenstein:

Within interactions-

Matt Backus:

So I'm going to do this at the interaction level. So it's true that some buyers might be switching back and forth, but across different observations.

Shane Greenstein:

Okay. That's pretty key because, yeah, if they're frustrated by lack of messaging on the mobile, you would just think they go their desktop.

Matt Backus:

Yep, that's right. But because they, so why is this positive for mobile buyers? Right? This is sellers sending the buyer a message, which the buyer can't see, right? Because they're on the mobile platform. Okay? So the seller doesn't know whether you're on a desktop and mobile. Right? So they might still get some messages. And by the way, it's not clear to us, we're not certain whether they couldn't see it at the time or not, we're going to treat them as if they couldn't.

Matt Backus:

And, econometric ally this turns out to be the way to get conservative estimates, right? If I do it the other way, just blow up my estimates by a factor of two. I'm happy if you want to do that. But we didn't do that from the payment. Okay. So the other thing to notice here is that adoption for desktop buyers on these interactions only goes up to about 6%. Right? So it's reasonably rare.

Matt Backus:

The downside of this is that it makes it really hard for me to identify stuff. The upside is that it means I'm not really worried about general equilibrium effects. Right? And if this were going to have as dramatic of an effect as I'm going to show you, it does, if there were higher adoption, then you would rightly criticize me for obvious general equilibrium effects, that I'm ignoring, but lowered adoption rates. That's going to be my get out of jail free card for general equilibrium. Yeah.

Williams:

So in the US let's say, which has had this, what's the percent of messages that people exchange?

Matt Backus:

Yeah, it gets up to about 10%. Yes. Higher interactions-

Williams:

[inaudible 00:51:50]

Matt Backus:

Yep.

Natalia Wright:

And that just based on the US pictures you showed us before, if we looked only at the top core tile and transaction asking prices, we would get a higher number two in Germany as well.

Matt Backus:

Yeah. You're more likely to send messages for more expensive stuff.

Natalia Wright:

So then, you, okay-

Matt Backus:

Yeah, just consistent with this being costly technology. Yeah. Okay. Here, I'm restricting attention to the case where you have buyer messages only, right? So now we're ignoring those seller messages to buyers who can't read them or I'm assuming, can't read them. And you get something pretty close to zero from the case, from the buyers on mobile. It's not exactly zero because if you've ever worked with data from a company, right, managing a large database is hard. So you never get those exact euros right? Stuff gets misassigned who knows what's happening. There's no documentation. The person created the table. I quit the company a long time ago because it's the deck world. So yeah.

Williams:

What's the volume of transactions again here that you're looking at? How many are we talking about?

Matt Backus:

So that's a great question. Let me get to my summary sets up as it's slipping. I think, it's about 3.4 million is the set that I'm going to use for four weeks before four weeks, the change. Great. Okay. So 3.3 million interactions. So I've split this. I've done this table a little bit strangely, by listings, by sellers, by buyers and by interactions. So by listings. We have a bunch in our data set, average asking price about a hundred dollars, and about 56% of those listings are going to sell. Okay?

Matt Backus:

At the seller level and the buyer level, the only thing I want you to know is that there's a long tail of very active buyers and very active sellers, but on average those buyers are much less active. So you have a lot of pseudo professional sellers. The modal buyer shows up once and then leaves the site. Okay? So I'm going to think about sellers as longer on the players. Buyers is short run players that'll matter later, not yet.

Matt Backus:

And then here's the interaction level data, right? Different from the listings, the listing might have multiple interactions. They get about a little bit less than two offers in their lifetime. 44% of the time it's successful and about half the time the buyer is on desktop. Okay. And mechanically half my sample is going to be after the period because I'm doing four weeks before and four weeks after.

Matt Backus:

Okay. I have a bunch of co-variants that I will throw in sometimes on time trends. Sometimes I'll take it out, you'll see why. The asking price category by condition, fixed effects. Because I think used means different things in different categories, day, week, fixed effects. There was some anecdotal evidence that in the UK it really matters whether it's rainy or not per activity on the site.

Matt Backus:

So I threw in weather and Frankfurt, it doesn't really matter. And then there's a ton of holidays in Germany in may. So I'll throw in fixed effects because obviously, four weeks after it's going to think... Okay, sorry. So first let me show you my identification strategy. See I'm doing it perfect. So there's two ways to think about identification here. What I couldn't do right?

Matt Backus:

If I were a product manager at eBay, what I wouldn't do is a pre post estimator with no standard errors. Right? Just take the mean before, I mean after, and what we're going to do based on that desktop mobile strategy is we're going to do a diff and diff, right? By comparing pre-post in mobile and desktop using the mobile group who never got the opportunity to send messages or couldn't see messages as a control group to deal with secular trends.

Matt Backus:

Okay? Cool. So let me just show you this in pictures to start off. This is the pre post for the desktop users. Okay. Here, I've excluded the time trend and I'll tell you why in a second. But so on the X axis is time on the Y access is the fraction of interactions that successfully ended in a transaction. You see it's really noisy, right? And there's this tiny half a percentage point bump.

Matt Backus:

That shouldn't be surprising that it's small, right? Because there's only a 6% adoption and this is site-wide success. When I add my co-variate in excluding the time trend, it gets a lot less noisy and we don't kill the bum. Why have I excluded the time trend? If I included it here, I'll show you in a second. The regression results I included here. Notice in the post period there's this trend.

Matt Backus:

So it turns out if I include the time trend, I'm going to kill my result. Okay. So here and this is part of why the diff and diff is going to be important for believing the results. Oh, sorry, before I do that, let me just show you the mobile, right? Mobile pre no difference, sorry. A mobile raw data, no difference. Mobile with controls, no difference. Okay? So this is what I'm going to get out of the mobile sample.

Matt Backus:

Is a little bit of credibility for how I'm treating this trend in the post period. That's really what I'm looking for. Okay, so let me show you the table of results. Before I do that, there's just a quick aside on issues of power. So it sounds like, when I say I have 3.4 million observations or 3.3 million observations, it sounds like a big data exercise, which is what you might expect from somebody who has internal access to eBay data.

Matt Backus:

This is not really a big data exercise, right? So we have 6% adoption. So the effects on the whole sample are going to be much smaller. And you can do this as a power calculation in the same way you would if you were designing medical experiments, right? So we'll say there's a binomial with a success rate equal to the sample average. The effect is about half a percentage point.

Matt Backus:

Okay. Say we want a type one area to 0.05 and say we... And here I'm just picking something arbitrary. You want a 50% power, then what you need is 3.41 million experiments. I have about 3.3 million observations. In other words, this is not big data. This is on the edge of maybe being able to detect something. I'm putting this up because at some point, pretty soon you're reasonably going to start asking me did you cut the data this way or that way?

Matt Backus:

And I'm telling you that if I start cutting the data extensively, I'm not going to see anything. Right? My standard errors are going to get big. I'm going to have a multiple comparison problem and it's going to look like a lot of garbage. So this is the defensive point. Okay.

Shane Greenstein:

Isn't that also an argument for a different diffing with as many different kinds of diffs as you can find?

Matt Backus:

Yeah.

Shane Greenstein:

I mean, you ruled out using other countries as your diff-

Matt Backus:

Yeah it's really bad.

Shane Greenstein:

Yeah, exactly. Immediately I know my responses. Well then go Austria, go to France, go to the US go to anywhere else because your time trend potentially might be identified, by seeing the same trend in some other country or something. [inaudible 00:58:34] isn't a reaction to have to this slide you just put up with-

Matt Backus:

I mean, we did look at other countries and it's a really bad comparison. So I could show you that there's not, we could go back and just show you that there's not this trend in other countries.

Shane Greenstein:

Yeah-

Matt Backus:

That'd be convincing. We can do that's easy-

Shane Greenstein:

For example, just because it's so small.

Matt Backus:

Yeah. That's a reasonable question. We can just do that as a separate set of plots.

Shane Greenstein:

Yeah.

Matt Backus:

There's eight other countries in Europe. We don't see this uptick during this period. I think that would be useful. Yeah. I like that a lot I should maybe write it down. Okay. One other thing, that some people only familiar with, so I'm not, so when I do the diff and diff, what I'm going to get, because not everybody uses the feature, right? All I'm getting is an intent to treat estimate, right?

Matt Backus:

It's like I gave you the pills. You didn't necessarily take them, I'm giving you the option to send messages, but you didn't necessarily use it. So this is what we call an intent to treat estimate. So that's what I'm estimating with the diff and diff, what you might want to do is ask, well if only 6% of people adopted, how should we blow that up to think about the treatment effect on the treatment and the way you do that as you run an IV.

Matt Backus:

Where you use a dummy for being in the post set who did send a message, as an instrument. The reason I'm highlighting this is because when I show you IV results, it's not introducing, typically we say we're using an IV. It's like we've introduced some new source of variation. It's not, it's the same source of variation. All it's doing is rescaling the estimate. Okay. So this is the table of results that corresponds to what I showed you graphically.

Matt Backus:

So first let's start with the pre post estimates. Just for desktop users. I don't use controls, right? I checked boxes empty, then I get about a half a percentage point of fact. If I introduced controls with that time trend, it dies, right? And that upward trend in the post period. Okay. So now what I want to do is I want to say, well let's use mobile to pull out that time trend, right? Because you saw on those pictures, mobile's flat.

Matt Backus:

So if I introduce mobile now the pre pulse for just mobile, I get a zero. This is the flatness of that mobile curve right. And then when I do diff and diff, now it's the interaction of post and desktop that I'm interested in for my estimate, it survives without controls and now instead of losing it with the inclusion of a time trend as we did before, we keep it. And that's because mobile is identifying the time trend.

Matt Backus:

So what is it doing? Then it's saying, well there must be some increasing treatment effect in the post period. Okay, I'll show you that, then when I reformulate this as an IV to say not what is the effect, the intent to treat for the marketplace at large, but instead for that 6% of people that were actually sending messages, what was the effect for them? I get this 0.0744 so let's think about magnitudes.

Matt Backus:

This is about a half a percentage point for the intent to treat for the marketplace at large. This is a seven percentage point effect for the people who actually send messages, right? So the baseline was about

44% so this is a 14% increase in the likelihood that we transact. And coming back to an earlier question, right, do we want these people to transact? Maybe not, right?

Matt Backus:

Maybe Shane likes his guitar more than I do, right? So some fraction, those failures are not failures at all. Right? But an efficient outcome and viewed that way. This is actually a conservative estimate for how significant the effects on the bargaining marketplace are. Okay? So these effects to me are astronomical. Now, we interpret a treatment effect on the treated, it might be that the people who get the most out of sending messages are the ones who send them as you would expect in a rational model.

Matt Backus:

So this is going to be an overstatement, most likely. Still pretty big, I think. Okay. And by the way, we're working with eBay right now to try to get them to roll out an individual. I remember there was that extra box add a message to a seller. Right? It was a perfect setting for an experiment to nudge people by removing that extra click. And we're working on that experiment to then get at some of the heterogeneity and these effects. And hopefully one day [inaudible 01:02:41]. I've been trying for years.

Chris:

Is there a lot seller level variation in the adoption of the some like... Some people would probably [inaudible 01:02:49] Kathy's, right?

Matt Backus:

Yeah. Some people are trying the Cathy's, some people send standardized messages. Yeah. A lot of seller heterogeneity in this.

Chris:

So if people pick this out over time?

Matt Backus:

No. Your adoption rate was pretty flat.

Chris:

[inaudible 01:03:09] Did we realize to keep working?

Matt Backus:

That I don't know. That's harder. The adoption rate was pretty flat. I mean, maybe you saw, let me see if I can-

Williams:

I thought you said you saw a trend over time of increasing rates. Right?

Matt Backus:

So maybe you see a slight trend here, right? That's pretty small, but you could probably foresee when I'm going to show you soon, which is the treatment effect is definitely increasing, but it's not that they're

using it this much more, right? That was relatively flat. It's that some reason it seems to be having a larger effect on the marketplace.

Williams:

They get more comfortable using it or something.

Matt Backus:

Yeah. So that's right. So if you were incredibly generous, you would say, Matt, your identification strategy is fine. This must be an increasing treatment effect. Right? And then we would cook up a story like maybe they're learning how to use messages or sending different words, right? That's the generous interpretation.

Matt Backus:

The less generous interpretation is, "Hey Matt, you're identifying off the 60% compliance. So it doesn't take a lot of misspecification for you to get an effect and a statistically significant one." Which is why it's going to be important. Then for me to put some meat on the bones of this convergence story. Yeah. Okay, so I showed you that. So now let me do a little bit of about Sines and then we'll dig into that convergence story.

Matt Backus:

So, this is actually what I was going to go in the direction of what I was just telling you. So the next thing that I want to do is think about, week specific events, right? So this is coming out of a paper by Dave [inaudible 01:04:38]. The idea is we're going to interact, the dummy for treatment with weeks specific fixed effects. And we do this for two reasons. One because this allows you to take all the weeks before the treatment and do an F test on those, right?

Matt Backus:

What does that testing that's saying are those effects always zero, right? And that's important because if you get a causal effect before the treatment, then time and space have failed us in some way that we haven't don't fully understand, right? Or you're really miss specified, right? So when people think about this as a check of the assumptions that underlie this diff and diff identification strategies, I'm going to do that.

Matt Backus:

But two having these weeks specific fixed effects is going to let me capture then that time trend. So here's what I get. By the way, you should, this is exactly what you should have been expecting right from that scatter plot. Right? So you saw the increasing trend in the post period. So here's the pre period I get zeros, thankfully. Right? I wouldn't be giving this talk if I hadn't gotten zeros here.

Matt Backus:

Because then you would know that I messed up the identification. Everything right is set to zero for the week before the treatment. And then in the period after the treatment we see this increasing trend in the treatment effect. Yeah.

Alex:

Why do you expect increasing trends? Why isn't there something jump? Is it because there is no learning story [inaudible 01:05:57].

Matt Backus:

Well, yeah. So I mean I didn't know what to expect. This is the data speaking. You're right that I... I think approaching this, I expected a sudden jump, especially after I saw that immediate adoption. Right? If this has an effect on communication, the simplest model you treat me, I should be treated. Right? And so when we see this now we have to go fishing for stories again. Right?

Alex:

You could build a learning story from this-

Eric Markin:

I guess you could-

Natalia Wright:

I mean especially if you got sellers who are selling all the time buyers who are making the same [inaudible 01:06:33] mistakes all the time.

Matt Backus:

Yep.

Natalia Wright:

Right?

Matt Backus:

Yeah-

Natalia Wright:

[inaudible 01:06:36]-

Matt Backus:

Yeah, no, I mean I'm going to try to make lemonade out of this, right? So what are economists really, famously bad at we give you equilibrium but we don't tell you how people get there, right? We don't talk about convergence to equilibrium and I'm not going to tell you how they get there, but what I am going to argue is that this is a story about convergence, right? This is that, you know what, you gave people a Cheap talks again, let's think about the complexity is equilibrium right?

Matt Backus:

Because it's a general message based. You could type anything into that box. You can type the elephant is sitting down. It's like what does that mean? Right? So the message... And the meaning is a feature of the equilibrium. And so you might imagine that this is actually a hard space to learn or to find an equilibrium in.

Alex:

[inaudible 01:07:19] the buyers are ready come only once or twice at best. So, There's not learning by my house right?

Matt Backus:

Yep. That's true. That's going to be really important. Yeah, that's exactly right.

Speaker 16:

Sir?

Matt Backus:

Yeah.

Speaker 16:

[inaudible 01:07:34] they only learn [inaudible 01:07:34] is this center of learning how to be strong. [crosstalk 01:07:39]. I was responding offended because my guitar is [inaudible 01:07:47]. But then I learned how to communicate with people in a way that I actually ended up with the say, I don't know.

Matt Backus:

That would be a compelling story if somebody told it-

Chris:

I have a slightly different reaction with it's 0.01 which is much higher than your estimate, which was taking the average of those four facts. Which suggest then that your prior estimate is a gross under estimate of the net benefit. And I would have been very curious to have gone out many weeks. Yeah, exactly.

Matt Backus:

Yeah. Good idea. Yup. So what are we seeing in the effect even like jumps up a little bit higher than we expect for a couple of weeks and then stabilize right about at that point if I know that by the way, because this was my estimating sample so it really is a little bit of an accident, right? That the mean over this was equal to the long run mean it's reassuring for if you want it to extrapolate from my estimates that this really does seem to stick as the long run mean but you're right then it was an accident, right? Because we had these do higher ones. Yeah.

Mike Wills:

So I would think that there are other bits of data that could help to inform the story. One is what happens to the variants and transaction prices more. Do you see people accept offers that look like offers that probably wouldn't have been accepted before the message?

Matt Backus:

Yeah.

Mike Wills:

Meaning like seller[inaudible 01:09:16] now determines that once they see a message they actually accept something that you didn't see them accepting for a similar item before.

Matt Backus:

I'm a little bit nervous about the variants just because getting it second moments is really going to push me on power. And, and the second case, it must be mechanically true, right? Because the average acceptance rate is increasing. And, the question is where is this heterogeneity, which is something we're going to do-

Mike Wills:

Yes. Whether the support accepted prices changes or if it's a changes in what's accepted within the support.

Matt Backus:

So the way that I'm going to think about getting that it's going to be, by cutting the data and ask them, do we see the changes in the low price bins? What do we see the changes in the high price bins? Is that getting era tuition or something different?

Mike Wills:

I didn't actually know what to expect, but I love it.

Matt Backus:

So I'm going to do this by cuts in the data. If we can find a more clever way to do it, I would really like to, because, as I said, when I start cutting the data to get at these heterogeneous effects, then I'm going to lose power and you're going to see it. It turns out pretty bad. But let's get there. So, okay, so again, the generous story here is that this is convergence to an equilibrium.

Matt Backus:

But again, it would be totally reasonable for you to say your compliance rate is really low. Right? So that's too much. Probably this is just some misspecification that's blown up by the low compliance. Right? Okay. So what can I do to convince you? One thing I can do is I can throw in seller effects because eBay.de is great and I have individual IDs and I throw in seller fixed effects and everything survives. So that's some good news. The other thing I can do is look at compositional changes. So this is a second answer to the-

Natalia Wright:

[inaudible 01:10:58] that's like seller experience, right? It's like a time varying seller thing, not, because like, yeah, but this side of the table was saying was about seller learning-

Matt Backus:

Yeah. I'm going to get there. Although your question is still going to be the right question to ask when I do that-

Natalia Wright:

Okay. But are we [inaudible 01:11:17]-

Matt Backus:

Yep that will be the right question.

Natalia Wright:

Yeah.

Matt Backus:

Which we don't do yet, but should.

Natalia Wright:

Yes.

Matt Backus:

Yeah. This is something much simpler, which is just asking, take the observables right and ask is the predicted probability of success for the things that we're bargaining over changing? Or another way of asking this is the set of stuff that we're negotiating over changing and we don't see any significant effects. This would have been more of a general equilibrium concern. Okay? So the compositional changes don't seem large here.

Matt Backus:

And none of it is statistically significant. We tried playing with the sample window a little bit and you could probably... So here on the bottom now this is all of these are all individual regressions in the bottom or the IV estimates. You could probably have predicted this based on just the time path of those weeks specific effects, right? It's small in the beginning it gets largest when you include those high weeks, and then it stabilizes a little bit lower.

Matt Backus:

Okay. So now, we started to get at this question and mechanisms a little bit cutting the data. And I'm pretty skeptical of this portion of the paper because we're reducing power. We're introducing a multiple comparison problem, right? We're just searching for effects now. So we did it two ways. By price buckets and by major categories. Yeah, I use strong language if we cut garbage. So let me at least show you the unsuccessful results.

Matt Backus:

So this is my price spins. So zero to 50, 50 to 150, and we're literally just cutting the sample and rerunning those regressions. These are the IV regressions. And one thing that was a little bit surprising, but I hesitate to put too much on it, is it seems the largest effects are for cheap stuff. Right? But remember, I already told you that we send messages most and more expensive stuff.

Matt Backus:

So this was a little bit surprising, but I'm not totally sure how to interpret it. Right? It might just be that there's more surplus at stake here. And so it still might be higher returns to send messages for expensive stuff.

Gram:

[inaudible 01:13:17] did the sellers change your asking price behavior? Now they know they didn't-

Matt Backus:

Yeah. This is the general equilibrium story and if they did, it would show up in this compositional picture. So, I don't think so. Yeah. But that's my best answer to that. I don't know. One thing I could do, which we haven't done is we could try to isolate the sellers who have some experience with messaging and see whether that compositional picture holds up for them and that might push it a little bit further in the spirit of your last question.

Shane Greenstein:

Probability of breakdown higher or lower as you get to the more expensive items?

Matt Backus:

Much higher.

Shane Greenstein:

Much higher?

Matt Backus:

Yeah. The probability we transact is pretty high for cheap stuff and pretty low for expensive stuff.

Shane Greenstein:

So back to the value would generate possibly more messaging even if it's-

Matt Backus:

Yeah, that's right. That's why I hesitate. Even though these magnitudes are smaller-

Shane Greenstein:

Yeah, all right. I was following, I was a step behind you. Okay. Gotcha.

Matt Backus:

Yeah, maybe if we did it in log percent or something, it would, but yeah, I'm not sure what exactly the right one is. So we just put up the raw, then we did it by categories. So these are categories for which we have at least a hundred thousand observations. We've got results all over the map. These are not corrected for the multiple comparison problem, these standard errors, so you shouldn't take them too seriously.

Matt Backus:

We tried to read the tea leaves for a little bit, but I have no idea of why, for instance, PC and video games, we see a negative effect, Y for mobile phones we see a strong positive effect. This was a bit of a dud for us. We were hoping to see per your question earlier, right? Maybe in the massive, incomplete information categories. We see more and I'm just not seeing it.

Matt Backus:

The other thing to remember by the way is whatever your representative idea of an item for these categories is it's going to be weirdly selected when you think about best offer, right? Where's where my DVDs going to be? TV, video and audio. It's not the retraining DVDs because you're less likely to use best offer for that, but maybe a led Zepplin collector set. So that selection is going to sort of distort what actually gets put in here.

Matt Backus:

It makes it really hard to interpret category fixed effects in general for selected eBay stuff. Okay. I don't like the slides so I'm going to skip. One other thing we did is we looked at other outcomes, right? We've been putting the probability of a transaction or the dummy for the transaction on the left hand side. Let's look at the number of offers, agreed price and so on. We don't see any significant effects. I'm going to focus on this bottom row.

Matt Backus:

That's the intent. Sorry. The treatment effect on the treated estimates. Here I'm giving you the, intent to treat estimates in the first two rows. We don't see any significant effects on the number of offers or even the first offer. We do see a significant effect, and again, maybe there's a multiple comparison problem here, we throwing stuff at the wall, but we do see a negative effect on the agreed price.

Matt Backus:

And so we decided to break that open a little bit by splitting the sample even more conditioning on agreement for whether the buyer's first offer was agreed to or a different sample where the seller's counter was agreed to and a different sample where the buyer's counter to the seller's captor was agreed to. So in other words, at what point in the negotiation did we conclude and what we get, are these alternating signs, which-

Gram:

And no significance-

Matt Backus:

They're not significant, that's right. The alternating signs suggest to me that, basically this is a story that whoever's sending the message is using it to their advantage right? The buyer offers accepted. Then they got a cheaper price, the seller's counter is accepted and they got a higher price. So messages are being used to push the price around, which is why then maybe we get a negative effect here, right?

Matt Backus:

Because buyers always make the first offer and those are often accepted. So let's just compositional mechanical feature of the platform. Okay. So now with my remaining time and we're just seeing what you have. Perfect. I'm going to talk a little bit about those messages. So in the post period, we have four weeks. I have about 10 minutes. Yup. Perfect. We have four weeks of messages that are sent on ebay.de that's about 250,000 messages.

Matt Backus:

So we start with cleaning and by the way, on eBay.de it's not like the US where there's an expectation you're going to use English. Only about 80% of the messages are in German. So we had to do some

language detection before we could even start de stemming words. So removal of numbers and symbols and stop words. And then limitizing just means you're chopping off the tense stuff. All the letters that just make up variants of a word. Okay.

Matt Backus:

I want to emphasize this is totally descriptive. I have no experimental variation here. I can't tell you any causal effects. The neat thing is that I'm going to be able to match this convergence to equilibrium. So several points I've been saying, this is the generous interpretation of my findings.

Matt Backus:

What I want to do is show you something else in the data that lends credence to this interpretation to try to convince you that it's not just misspecification driving identification off of a tiny compliance. Okay, here's some summary of that data. We throw a little bit of stuff at the wall here. We tried to compute the sentiment so they're... In text analysis there's this stuff called sentiment analysis where you just ask how nice are the words that people are using.

Matt Backus:

And it's literally there'll be a dictionary with numbers scores, right? Thanks is going to get a positive score or jerk is going to get a negative score. And then you just add up the score across all the words. It seems buyers are a little bit nicer than sellers. This might be mechanical though, right? Because sellers are responding to an offer. German makes up about 84% of the sample, among buyers, right?

Matt Backus:

Among sellers, you see a little bit more use of English, other languages. It's pretty similar. Here I can show you the most frequent stems that show up. Hello, a lot, a little bit of talk of shipping by the way. And this question of whether people negotiate over shipping with the messages. Technically they're not supposed to in the instructions for best offer. Clearly they are at least talking about it if not to go shooting on shipping. A little bit more nice stuff on the buyer side, by the way. MFG is, you might have to help me from the gruesome?

Natalia Wright:

It's [inaudible 01:19:18] yes. Its like a [inaudible 01:19:20].

Matt Backus:

Yep,

Natalia Wright:

It means like you be like [inaudible 01:19:23]. It's like a casual-

Matt Backus:

This baffled us for so long. Okay. So anyway, buyers say a lot of nice stuff, sellers are talking to, still some hesitant language. Unfortunately that might be the sentiment difference.

Anne:

What did you do with it? Is that, yeah. So some of those words mean different things in German and English. So the word like also, in German that word literally means therefore.

Matt Backus:

Yep.

Anne:

So, I'm just confused about-

Matt Backus:

What have I done? So we de stemmed in German and then to make that table, I put it through Google translate. So hopefully it picked up on that.

Anne:

Okay, I see.

Matt Backus:

Yeah. Everything here I've just smashed through Google translate because I-

Anne:

[inaudible 01:20:06].

Matt Backus:

Yup. Okay. So now-

David:

Do we own the English or own the German. Which one is the analysis?

Matt Backus:

The analysis is done in German.

David:

Okay.

Matt Backus:

Yep. Which we made things a little bit harder because there's not as many sentiment analysis, dictionaries and more tools in English. And also I have no idea. In some ways it's better to not know what your code is doing. So what do I want to do with this? Here's the payoff to doing this text analysis stuff. Do you remember the story about convergence to an equilibrium? The question is can we see it in the messages people are sending.

Matt Backus:

So what I'm going to do is I'm going to take a Bag-of-words approach and by this is a learning experience for me in natural language processing. It was shocking to me how primitive the state of the art, natural

language processing is. Bag-of-words is like state of the art probably until what? 2013, 2014 when we saw a word to back and then later [derk to back 01:21:00] But this is like real NLP stuff, which blew my mind.

Matt Backus:

Okay, so we're going to do this separately for buyers and sellers and what I'm going to do is take all of the messages that were sent by a buyer in week one and squish them all together in one document. We to squish them all together in one document. So now I have a big long document for buyers for weeks one through 10 and for sellers, for weeks one through 10. And then I'm going to do Bag-of-words.

Matt Backus:

And let me just see if I can, I'm going to try to click on something just to explain what [Bag-of-words 01:21:31] is. All it is. Literally Bag-of-words is just converting a document into a vector, right? Where each element of the vector is a counting measure over the support of words that are ever used in the joint of any of these documents. So a practical example, your document one, thanks for your offer, we [inaudible 01:21:50], we get rid of stop words and that becomes thanks, offer or thank offer probably I should've picked sound side, right?

Matt Backus:

And so what does that, I'm going to say element one of the vector is thanks element to his offer and on the three is going to come from some other documents. So you get one one zero and if I had said thanks, thanks for your offer, it would be two one zero okay. Document two. This is my final offer. Now we have this word final, which didn't show up here. That's going to be in the third position offer was in the second position. Right? And so it's zero one one.

Matt Backus:

So we're converting, these documents and you can do this for arbitrarily long documents, right into vectors. And once we're in vector space, now we have all kinds of similarity measures, right? That's why we're doing this. That's why Bag-of-words was useful. Okay, cool. So now I have 20 vectors, 10 vectors for buyers, 10 vectors for sellers. And what I'm going to do is take the co-sign distance. This is a measure of the angle between two different vectors, right? In order to think about how similar are the sets of words that we're using.

Matt Backus:

I want to just a distance metric. We could have used lots of different distance distance metrics. Here's what I used. Okay. And the idea is that if we are converging to an equilibrium and that distance between say week one and two and week one and three, that should be larger than when we've converged. Right? Weeks five and six or seven and eight should be much more similar in co-sign districts than they are to the weeks where we were still converging. Does that make sense?

Matt Backus:

Okay. Alright. So you're the pictures.

Speaker 10:

[inaudible 01:23:23] it's funny. I would have tested something differently really. I would have tested a reduction in the variance of the words that get used over time as transactions get more standardized. So it's-

Matt Backus:

What's variance of the word mean?

Speaker 10:

Well, I mean I'm taking the actually from the patent literature which uses, that's something very similar to look at the uniqueness of words and patents as a way of measuring how novel the patent is. And so when you depart from commonly used words, you're more novel. Whereas here, the point is the transactions are becoming more regularized if it's a learning story and they're getting used to going back and forth. So the sellers who are guests are the ones that are learning here are paying attention to what words work. And so they'll all end up converging on this smaller set of phrases.

Matt Backus:

I like this idea because I think I can do it and I'm pretty positive it's going to be consistent.

Speaker 10:

You got the idea?

Matt Backus:

Yup. This is a fun idea. I hadn't thought of it and you should do this. I totally agree. Okay. So if the story is right, so are these? These are heat maps of that co-sign distance. Sorry, I thought [inaudible 01:24:32] to do zero is one. And so just shift everything up. Zero should be one, two should be three and so on. So what is this? This is the coastline distance between any pair of weeks, right?

Matt Backus:

So on the diagonal here, a thing compared to itself is identical. And so the darker the color, the more similar it is. So that's why it's perfectly black along the diagonal here, right? And then here we see the greatest distance between the vectors. And what is this? This is the difference between week one and week 10. Okay. So what should we expect? If there's convergence to an equilibrium on the far off diagonals, we should expect to see more distance.

Matt Backus:

And then if we fully converged, we should expect to see a block that is relatively more similar, which is symbolized in the heat map by a darker color. So here I've just replicated that treatment effects picture, right? So you can see that this roughly matches that convergence and a similarity of sets of words. Now what's really cool to me coming back to the question earlier, right?

Matt Backus:

Is well, buyers are one-time players so they can't be learning anything. And we do this for buyers, [inaudible 01:25:34] right it's the long run players who seem to be learning something, about what should be said or what words being in equilibrium. And then of course the next reasonable question is could we do this within buyer and we just have [inaudible 01:25:47]-

Speaker 16:

You care almost a little bit too about that which round of communication it is to like... You would think that the seller would learn how to respond to the first because if the buyers are always on average going to send the same first message, which is, I really want this guitar for my kid. Can you come down in price? The seller, they should learn how to respond best to that over time. It's like the first response from the seller or something, you know what I mean?

Matt Backus:

I mean I think it's going to be a little bit of authority, right? Because the buyers are just sending all sorts of crazy stuff. Buyers are one time players so they're not sending consistency-

Speaker 16:

Yeah, whatever that response is. To me it seems the seller just, I don't know, something about which of the seller messages.

Matt Backus:

Yeah. I mean your intuition that sellers are standardizing is right. Because when you look at the data, right? There are a bunch of messages that show up over and over and it's because the seller has decided this is my message and they copy and paste it. Right? And half of it is advertisements, half it is message. Yeah, actually the most populous message was there was an API that allowed sellers to upload and if you're thinking about designing an API, right?

Matt Backus:

It'd be really useful if there were a little grayed out text. Like, add your message here, promise this API, didn't add it as grayed out text it, prebuilt it as had your message here in German. And so that was the most commonly sent message. It was add your message here. I cut those out of the sample by the way, for the whole analysis. Okay-

Tony:

You should wrap up.

Matt Backus:

I should wrap up? Oh, yeah, sorry I'm out of time. Okay. So anyway, let's do that. So I never really worked out exactly what's changing. I've tried it looked sentiment analysis. Are they getting nicer or are they using particular words more? I just haven't figured it out. And so right now this still sits in this relatively high dimensional or to make sense of space. So, this communication facilities successful bargain and our results, absolutely.

Matt Backus:

And it looks like something where we can see that story of convergence. There's a lot of questions left open. What are they saying that affects outcomes? I can tell you what I can't prove yet, which is, I'm pretty sure that the effective messages are, to my dismay as an economist, not the rational, I'm doing this for a long time and I have a lot of information. I think it's the nice person stuff. Thanks for considering my offer. Have a nice day.

Matt Backus:

That seems to be where the biggest impacts are. I haven't written the table to prove that yet, but that's my intuition of where this paper is going. But, what is the mechanism by which this works? Understanding that and what's the optimal communication protocol. That's what eBay wants. And in general, coming back to this bigger picture agenda on bargaining, finding new sources of data so we can tackle these questions. All right. Thank you very much.