

Keynote Address

Joseph Farrell

The female mosquito will bite you and take your blood, but not much of it—perhaps less than 1/100th of a milliliter. For comparison, the Red Cross takes a pint, which is over 400 milliliters. I don't particularly begrudge the mosquito her tiny droplet of my blood. But I do resent the side effects, or in economic jargon, the transaction costs.

The World Health Organization estimates that malaria annually causes hundreds of millions of illnesses, and more than half a million deaths. Surely these transaction costs of mosquito bites vastly exceed the value of the blood actually taken. In payment instruments, of course, a major goal is to reduce transaction costs, and intuitively cutting those costs seems especially urgent when they are large relative to the transfer made. Transaction costs include time, and for small payments the time cost can dwarf the payment, as at some toll booths, or small cash register transactions with a long line.

The ratio of transaction cost to value transferred can be high if transaction costs are large, or if the value transferred is small. Getting that ratio down in the latter circumstance is, I think, what people mean by the problem of micropayments. Metaphorically, I'll call this the mosquito problem.

By that definition perhaps the micropayment problem can never be solved, because however low your transaction costs, they will be large compared to some payments you might like to make. Andrew Odlyzko has argued that micropayments will continue to disappoint. I suggest a more optimistic view. First, I think micropayments are a problem on which we can make progress, though not one we can solve: in fact, progress on micropayments is closely aligned with progress on payment instruments generally. Second, I argue that explicit micropayments in the sense of stand-alone small money transfers are not the only way to pay for small transactions: business creativity can work around the difficulties of doing that (as indeed Odlyzko has noted). However, qualifying my optimism, I point out that one widely used such work-around, namely advertising, raises privacy concerns.

We see a lot of experimentation, competition, and innovation in how we pay for things. Businesses experiment and optimize in search of more business, and customers choose the best offer open to them. This doesn't work perfectly, but it can work pretty well. Here, I will present a relatively optimistic view of the business side, and a mostly but not completely optimistic view of the consumer side.

If explicit micropayments systems are challenging, business arrangements can sometimes substitute. One arcane pleasure of economics is seeing how business creativity can handle such problems. I'll mention several broad strategies for the mosquito problem: that is, strategies to bring under control those pesky transaction costs for small transactions.

One strategy is the bundling of transactions to spread transaction cost. The ratio of transaction cost to value transferred is usually much lower when more value is transferred. Thus one can address the mosquito problem by bundling small transactions into bigger ones.

As one example, phone cards control pay-per-call transaction costs by bundling many calls. Following an upfront payment in money, you can draw on your credit with the phone card company. A privately created stored-value system, your account on the phone card, lets you skip some of the transaction cost of paying for a single phone call (though there is a transaction cost of using the phone card); and then you bundle those together in such a way that buying a phone card is a tolerably efficient transaction. Still on phone calls, a further step toward bundling is the prevalence of all-you-can-eat calling plans, saving on keeping track of calls made.

Newspapers are another example of this bundling strategy (as well as illustrating another technique below). A newspaper or a news magazine bundles dozens of news items into something valuable enough to justify newsstands and cash transactions; subscriptions of course further bundle these bundles over time.

Similarly, music albums bundle together different songs; one payment transaction cost is spread over multiple songs. But many consumers rather like to buy songs individually. To make unbundling small-value items work, you need to get the transaction cost down in nonbundling ways. One way that iTunes does that is by leveraging other payment instruments.

iTunes is trusted enough that many consumers are willing to have them store credit card information—this involves both trust about not abusing the information, and a password or other authentication system to control access to your iTunes account. This enables them to leverage off other payment mechanisms: when I buy a song on iTunes, they charge my credit card, which is paid by automatic payment from my bank account, which in turn is topped up by my employer.

But the central “mosquito control” strategy that I want to focus on hinges on economic complementarities, a key tool for indirectly processing payments.

In a two-sided market, or more generally with complements, cutting price on one side raises the demand curve on the other side, and a seller can profit with a higher price and/or more sales there. That added profit can be viewed as an indirect payment from the first customer, either in addition to, or in lieu of, explicit payment. An extreme form is to make one side free, which is particularly useful if transaction costs would be high on that side.

We see this in many contexts, notably newspapers and broadcasting, and in “free Internet content.” In those cases, the paying side has often been advertising. Ad support involves an indirect micropayment from consumer to content provider; willingness to be exposed to ads is a way for consumers to transfer value to advertisers and hence indirectly to the content provider. In many cases this value transfer is quite small; advertising thus serves as an indirect micropayment mechanism.

In that role, it has some real advantages. It works through payments and negotiations between commercial entities—no individual consumers fumbling in their pockets, and fewer security issues than many alternatives. People often say that “free” is a particularly convenient price; perhaps a better way to say it is that two-sided pricing with one side of the market being “free” may lower total transaction costs. And the transaction-cost gap between free and cheap might be even bigger in the Internet environment.

One example: newspapers. Some print publications don’t ask for money directly from the reader: they get all of their money—instead of the traditional roughly 80 percent of their money—from advertisers. As a result their newsstands don’t need an attendant or a coin box, and it’s quick and easy for a reader to grab a copy.

Another example is broadcasting. Traditionally both television and radio used the free (to the consumer) approach and were supported by ads. This was perhaps less of a choice and more a reflection of a constraint: they were perceived as non-excludable goods on the consumer side, so there was little possibility of implementing a subscription or direct-payment mechanism.

But once the possibility of charging consumers opened up, technologically or in regulatory terms, the question arose whether “free to the consumer” with advertising support was the best pricing model. Do you really want your visual entertainment interrupted multiple times per hour, as a way for you to contribute a few cents to the content creator? How much is your time worth? Perhaps this was not a very efficient payment mechanism after all.

Thus in television, and to some extent in radio, excludable forms of program distribution (principally cable and satellite), with subscriptions, were introduced. Ad financing was then partly supplanted by bundling-and-subscription, and partly supplemented by it: there are plenty of ads on cable. That slightly surprises me, as someone who finds repeated short ads distracting—but there seems to be a reasonable market test in there.

This brings us to the customer side. Often customer choice works pretty well, although sometimes it takes analysis and perhaps even a little faith to see this. For instance, sometimes a cash register transaction is delayed because a customer ahead of one in line is using a time-consuming coupon. I admit that I often vent about coupons and standing in line, but I also recognize it can lead us to an important economic point. The consumer can learn how long a particular supermarket's lines are. If a coupon scheme slows down the line, and the merchant doesn't add enough checkout counters, then some customers may no longer walk in the door. The merchant will take that into account. It's not perfect, but an only mildly over-optimistic view is that the merchant will weigh those effects pretty well. Similarly, the consumer facing part of newspaper or broadcast ads doesn't raise a lot of problems. Those ads are easy to ignore if you do not like them, and if you dislike them, you will likely know right away.

But, at this point in the development of our market institutions, customer choice is much less informed and not so reliably effective when it comes to privacy issues in targeted advertising. This was a substantial focus of the FTC's recent privacy framework report. From a payments system viewpoint, the issue is that ongoing changes in the ad-support micropayment model—changes that make it more effective, in some respects, by more tightly targeting ads — also weaken the presumption that consumers choose the best offer facing them. That presumption was a key to the market-mechanism argument that payment system evolution will lead to good outcomes. In other words, tighter targeting of internet ads may not fit the model that says the merchant—or the creator of content in this case—has an incentive to properly take into account any consumer harm. As a good deal of public policy discussion, including the recent FTC privacy framework, indicates, that remains a public policy concern.

In other words, what if the consumer cost of ads includes potential compromises to privacy and data security, rather than the simple annoyance of ads interrupting your programming? From an economic point of view, a key difference is the real risk that consumers are much less able to evaluate it and respond to it through their demand for the content.

Alternative (non-advertising) forms of micropayments could help. The “free content” versus privacy trade-off might be defused if we can make progress on paying through some alternative means—perhaps money, perhaps another indirect form of payment—in a way that reduces the pressure to serve targeted ads. We could get then a better market test of the efficiency and innocence of such ads; and if they are problematic, an alternative. If privacy concerns about targeted advertising can be assuaged through progress on other means for micropayments, perhaps the perceived trade-off between privacy and ready access to the long tail of Internet content can be relaxed. In other words, by working on payment systems, you may be protecting both privacy and the lively ecosystem of speech on the Internet.

ENDNOTES

¹<http://www.mosquitoworld.net/mosquitoFAQs.php>

²<http://www.redcrossblood.org/donating-blood/donation-faqs>

³See e.g., <http://www.who.int/mediacentre/factsheets/fs094/en/>; malaria is by no means the only mosquito-borne disease: see for instance <http://www.who.int/mediacentre/factsheets/fs117/en/>

⁴See Andrew Odlyzko, “The Case Against Micropayments,” <http://www.dtc.umn.edu/~odlyzko/doc/case.against.micropayments.pdf>

⁵An added benefit, for the phone card company and the user jointly, is often much better negotiating ability in dealing with the telecom company that supplies the actual telephone connection.

⁶This can of course go wrong if the user thinks he is in the all-you-can-eat zone and belatedly learns that he has ventured outside it. The FCC has expressed concern with this “bill shock” problem.

⁷For instance, I’m told that advertising rates, in ballpark terms, for a 30-second spot on over-the-air television are in the ballpark of \$5 to \$50 CPM: that is, per thousand viewers. That means somewhere between a half cent and a nickel per viewer exposed to the ad.

⁸See for instance Catherine Tucker, “The Economics of Advertising and Privacy,” working paper 2011, MIT.

General Discussion

Keynote Address

Mr. Tomasofofsky: That was very thought-provoking. I had a question on how you were talking about things. What would happen in the television subscription model where we are paying by watching commercials, if the technology comes along where we can press a button and fast forward through the commercials, what does that do to the rest of everything else?

Mr. Farrell: If you can very easily evade payment, then that can threaten the payment model. You might have thought that might contribute to the shift in payment mechanism toward subscription content—especially now remote controls make fast-forwarding even easier. But, after all these years of progress in fast-forwarding through ads, we still see quite a lot of ad financing of video content. Apparently, it was not a devastating problem.

As a side note, I've read that the initial response from the studios to the development of recording and potentially fast-forwarding devices was not so much they were worried about fast-forwarding. They were more worried about time shifting—you would watch the ad, but perhaps a day or a week later. They feared that some of their advertisers would worry about when—rather than whether—you would watch the ad. It puzzles me a bit, because my general sense is that with most television ads it would not matter that much if you viewed them a day or a week later.

Mr. Sullivan: Your mosquito problem is a very useful construct to understand this. I do think it focuses mostly on the transaction costs side. I wonder if you have had any thoughts about the risk side of this as well? When you look at, for example, the iTunes transaction, you buy a song, you get a 99 cent charge on your credit card. You buy another song and you have another 99 cent transaction fee. So they manage to deal with the mosquito problem with each of the 99 cent transactions.

What I am wondering is, Now they are getting that transaction charge for each one of those payments, why do they not just accumulate those over the course

of a month and then maybe put a \$10 charge on it? Maybe there is some risk side to that and there may be other issues that make the decisions to do this bundling more tied to some other aspects of risk.

Mr. Farrell: I do not know anything about Apple's relationship with its credit-card acquirers. In principle they might or might not have a per-transaction fee that would encourage them to accumulate charges. You could even imagine that they do accumulate charges but make it look otherwise to the end user. The key thing about the payment model is not what it does to Apple's relationship with its acquirers, but the effect on what the consumer has to do: he doesn't have to type in a credit card number and a billing address every time he wants to download a song. As I commented about advertising, the back-office operations between large commercial entities can often be made a lot more cost-efficient than can be done at the consumer end of things where there is a lot of literal or metaphorical fumbling in the wallet.

Mr. Morrison: I did have a question about your take and understanding of the pay-for-view ad, where there has been some very interesting innovation of where people are going to an ATM. Instead of paying an ATM fee, they watch an ad for a specific merchant or advertiser and, instead of paying that fee, the ATM pays for it itself. The other piece of that, there is a company called Jingit that does the same thing, except they actually pay people to watch ads. So it has a very directed focus on the ad and then they pay people to watch it.

Mr. Farrell: It might depend when the ad is selected. Some people might see a privacy problem if the ad is targeted to you once you have already identified yourself. I think that would be different if it were simply the ad that is shown at this ATM today and you press a button first to watch the ad, rather than pay the ATM fee.

Other than that, I would say it is an experimental business model. It might work, it might not. I recall back when long-distance telephone calls were a little bit expensive, there was a company that tried providing free long distance in return for listening to an ad on the telephone before they put you through. I don't think that particularly went anywhere, but it was an interesting experiment worth trying.

Ms. Hughes: The report that came out on Monday spends a good deal of time talking about not tracking. There is also a significant impulse coming out of the European Union to not tracking or even opting in. What are your thoughts about not tracking or opting in for an audience that is going to be doing payments providing in a variety of different spaces? It is not just Internet. It is lots of things.

Mr. Farrell: Probably I should stress again here that I am not speaking for the FTC. The Commission's report said what it believes. The report specifically said that it is applicable offline, as well as online. I think it also was not intended to get

at tracking of the kind that you might need for audit and security purposes. What the report is intended to do is to say the default should be “Do not track a consumer around the Internet, so as to be able to send more targeted ads or otherwise make money off the consumer, without appropriate consent or in the context of the transaction’s circumstances.” So it is a somewhat nuanced message, but basically the idea is this is a practice that is threatening to privacy and people should not do it, without proper consent.

Mr. Salmon: I completely agree with you about the positive externalities of bringing down transactions costs. Can you talk a little bit about what you have seen in terms of where transactions costs have gone over the past years or decades? It seems to me they actually have been going up rather than down. Also, can you tell us whether you think there is any real hope for them to be coming down in the future and, if so, where that hope might be coming from.

Mr. Farrell: I do not have a particularly privileged window on those questions. My sense is that transactions costs have been coming down, in part through innovation. Yet certainly dealing with a retailer online that you have not dealt with before or that has not hung on to your information involves a good deal of typing in information. That, in itself, takes longer and is more hassle than presenting a card, or for that matter cash, in a bricks-and-mortar establishment—where there has also been transaction-cost innovation. Of course, getting to the bricks-and-mortar establishment is part of the transactions cost too.

But I think the basic point remains: there is a really important agenda of getting those costs down further and that is what you payment industry guys, I hope, are doing.

Ms. Garner: Back in June 2011, Consumer Reports came out with an article that talked about fraud pretty significantly and mentioned the cost for the issuing community in the United States to move to chip and PIN was somewhere around \$2.85 billion, compared with \$2.4 billion in fraud losses. What implications does that have for the security of card payments going forward? In particular, as we look at mobile and we look to move some of these products we have today, a credit card as we know it, in the mobile wallet?

Mr. Farrell: On the mobile aspect, I will defer comment, because the FTC is having its workshop in the near future. In general terms, why not chip and PIN in the United States? I don’t know a lot about this; there have been conflicting accounts. As long as the industry, rather than final consumers, is bearing the costs for fraud, you would think they would have an incentive to introduce a fraud-prevention technology if it made sense. But there is always the possibility of incentive failures at various points along the line.