Barbara van Schewick
Assistant Professor of Law and (by courtesy) Electrical Engineering
Director, Center for Internet and Society
Stanford Law School

Opening Statement at the
Federal Communications Commission's Workshop on
Innovation, Investment and the Open Internet
in Cambridge, MA on January 13, 2010,
WC Docket No. 07-52, GN Docket No. 09-191

Good evening. My name is Barbara van Schewick. I'm an assistant professor at Stanford Law School and direct the Center for Internet and Society there. I also have a courtesy appointment at Stanford's electrical engineering department. I have a law degree and a PhD in computer science. For the past nine years, my research has focused on the relationship between Internet architecture, innovation and regulation. My book "Internet Architecture and Innovation" will be published by MIT Press this spring.

The Internet has created an enormous amount of application innovation. I'm here today to explain which factors made this possible.

I'll do this through three stories. Here is the first:

In the mid-nineties, a software engineer in Silicon Valley named Pierre Omidyar thought: "Wouldn't it be nice if we could buy and sell stuff over the Internet, using auctions to determine the price?" Most of his friends thought he was crazy: "Strangers will never buy from strangers online," they said. But Pierre Omidyar didn't care about that. He stayed home over Labor Day weekend 1995, wrote his software and put it online. Nine months later, so many people were using the platform that he decided to quit his day job and focus on it fulltime. The platform, renamed eBay, became a huge success, and today, more than 88 million people worldwide are using it to buy or sell things over the Internet.

Three aspects are important here:

Pierre Omidyar had an idea for an application. People thought it was crazy, but it didn't matter. In particular, it didn't matter whether network providers believed that "strangers will ever buy from strangers online." And that's because on the Internet, network providers don't have to do anything to enable new applications to run. This is a consequence of the Internet's architecture. The original Internet was based on a design principle called the end-to-end arguments. This design principle was first described by Jerome Saltzer, David Reed and David Clark who are here today. Following this design principle, the network was designed to be as general as possible in order to support a wide variety of applications with different needs. So when a new application comes along, the network doesn't have to be changed to allow the application to run. All the innovator has to do is write a program that runs on a computer attached to the Internet.

As a result, an innovator does not have to convince network providers that her application is useful, or will be commercially successful. The only person who needs to be convinced that this is a good idea is the innovator herself. This greatly increases the chance that innovative ideas will be realized.

Second: When the application has been written, the network does not need to be changed before the application can run on the network. If you want to use it, you install it on your computer. That's it. The only person who needs to be convinced that this application may be useful is the person who actually wants to use it. This greatly increases the chances that people can actually use the new application.

Third: In this architecture, it doesn't cost a lot to develop new applications. You need access to a computer, be able to program, and time to actually write the program. This greatly increases the number and type of people who can develop new applications. Like Pierre Omidyar, you don't have to be an employee of a firm or have outside funding to realize your idea for an application. Because the biggest investment is often the design and programming of the application itself, innovators can develop an application in their free time or as a side project. Under these conditions, an application does not have to produce a profit in the future to cover the costs of developing it. Instead, a wide range of benefits may be sufficient to cover the development costs.

Thus, the architecture allows innovators with a wide range of motivations and funding models to develop applications.

So, three aspects:

Innovators independently decide whether to realize innovative ideas. They do not need support, or "permission" from network providers in order to innovate.

Users independently decide which applications they want to use.

The low costs of application innovation enable a very large and diverse group of people to develop new applications.

That was the first story. Here is the second:

In 2002, two European entrepreneurs named Niklas Zennström and Janus Friis thought: "Wouldn't it be nice if we could use peer-to-peer software to make phone calls over the Internet?" At the time, most network engineers didn't think this was possible. They thought that Internet telephony would require special treatment from the network (something we call "Quality of Service"). Network providers weren't really interested in pursuing the technology because it was a huge threat to their business model. But Zennström and Friis didn't care about all this. They went ahead, developed their software, the software became Skype, and today, more than 500 million people worldwide are using Skype to place phone calls over the Internet.

Again, we have two entrepreneurs who had an innovative idea for an application. Network engineers didn't think it would work, but it didn't matter. Nothing new so far.

The application constituted a huge threat to network providers' business models, but it didn't matter. And for Zennström and Friis, it didn't matter because there was nothing network providers could do about it. And there was nothing network providers could do about this, because the Internet's architecture prevented them from interfering with the applications and content on their networks. As I said, the Internet was based on the end-to-end arguments. As a consequence of this design, the network couldn't distinguish between different applications and content (it was "application-blind"), and as a result, network providers couldn't control the applications and content on their network.

Today, that's different. Today, sophisticated technology is available that enables network providers to identify the applications and content on their network and control their execution.

Thus, the original Internet was application-blind, today's Internet is not. Does it matter?

Imagine you have this great idea for a video platform that will revolutionize the way people watch TV. Once they have used your application, they will never want to go back to cable again. Of course there are risks. The technology may not work. Users may not like your product. Your business model may be wrong. But in the application-blind network, you know that you will get a fair chance in the market place. You will be able to compete with other applications on the merits.

In today's network, cable providers may squash you. The network can turn against you any time and block your application or slow it down. There are many reasons why network providers may want to do so. Maybe your application competes with theirs; maybe they just want a share of your profits. Maybe they don't like your content, or your application is slowed down to manage bandwidth. Whatever the network provider's reason, if your application gets blocked, your project fails, and you won't be able to reap its benefits. And accounting for this possibility, you (or potential investors) may decide not to pursue your idea.

Third story:

When YouTube came on the market, it competed with Google Video. You Tube was better. It won.

In today's Internet, things might have been different. In an application-aware network, the network provider can ask applications to pay an access fee. There are many ways in which it could do so, and all of them would be bad for application innovation.

Let's focus on one possibility: When YouTube came on the market, network providers might have said: "Google, you are big. You have lots of money. Why don't you give us some of this money, and we will give Google Video better transport." Imagine Google pays. Suddenly, Google Video is so much better. Not because it's the better product, but because Google is rich, and Google was able to strike a deal. In such a world, network providers get to decide who is

successful, by deciding who gets a deal. Suddenly network providers, not users, get to pick winners and losers on the Internet.

Three stories, different factors. How do changes in these factors affect the amount and type of application innovation?

Some changes may affect the benefits and costs of innovation. An innovator decides to innovate if the benefits (broadly defined) are larger than the costs. Increase the costs or reduce the expected benefits (for example, through access charges or discrimination), and some innovations may not be justified any more.

Some changes may affect the size or diversity of the innovator pool. Others may let network providers, not users, choose which applications will be successful and how the network can be used. For example, access charges may reduce the profits of all affected application developers, but they may hit certain types of innovators (for example, those with no or little outside funding) particularly hard.

Why are these things important?

If there is uncertainty (e.g., about technology or user needs) or if user needs are very heterogeneous, a larger and more diverse group of innovators will produce more and better applications than a smaller, less diverse group of innovators, and that innovation will better meets user needs.

What's the intuition here?

If there is uncertainty, nobody really knows in advance which applications will work, or which applications will be successful. Under these circumstances, economic theory suggests that it is best to try out many different ideas, and see what happens. Some will succeed, some will fail, but trying is the only way to find out. And because different people will have different ideas and different views of the world, more and more diverse people will have more and more diverse ideas. Tim Berners-Lee looked at the Internet and saw a giant web of shared information; Pierre Omidyar saw an online marketplace, and Jeff Bezos saw an online bookstore. A larger and more diverse group of potential innovators will also realize more of the ideas that are known. For

example, start-ups may have an incentive to realize ideas that established firms wouldn't pursue. Users have an incentive to create applications that manufacturers won't produce.

By contrast, fewer innovators, or less diverse innovators, will try fewer things, leaving valuable ideas on the table.

But widespread experimentation is only part of the mechanism that produces innovation under uncertainty. The second is: Who gets to decide which applications become successful? Users or network providers? Does it make a difference?

I argue it does – because users and network providers will choose different applications. There are two reasons for this:

First, users and network providers use different criteria when choosing which applications will be successful. Users choose the applications that best meet their needs. That's easy. Network providers may use different criteria: "Does this application compete with my own application? Does it create a lot of bandwidth? Does my preferred vendor offer network management tools that happen to block this application?" Consider Skype. Many mobile providers in Europe do not allow their users to use Skype over the mobile Internet. If you look at user forums, you will see that users don't like this. They want to use Skype on their cell phone. But if users use Skype, they don't make as many traditional cell phone calls, and voice revenue shrinks. So network providers make a decision that's different from what users would choose.

And second, even in those cases where network providers would like to choose the applications that users want, they don't necessarily know what that is. That's the uncertainty I talked about earlier. In many cases, nobody knows whether an application will be successful until users actually try it. Network providers cannot replace this.

Beyond innovation, user choice is also important if we want the Internet to provide the maximum value for society – but that's another story.

Thus, if network providers pick winners and losers on the Internet, if they decide how users can use the network, users may end up with applications that they would not have chosen, and may be forced to use the Internet in a way that does not create the value it could.

In sum, there are a number of aspects that foster application innovation:

Innovators independently choose which applications they want to pursue; users independently choose which applications they want to use. The application-blindness of the network ensures that the network provider cannot interfere with these choices, that it cannot distort competition among applications or reduce application developers' profits through access charges. Finally, the low costs of innovation not only make many more applications worth pursuing, but also allow a large and diverse group of people to become innovators, which in turn increases the overall amount and quality of innovation.

But why do we care so much about application innovation? Why should policy makers care about it?

I have a longer answer to this question. It explains how application innovation contributes to economic growth and how it creates value for society in all areas of society. But my time is almost up. Therefore, let me just say this:

Did you ever try to explain to your partner's grandmother why she should get the Internet? I did. Although I'm a computer scientist, I didn't say: "Grandma, you have to get the Internet! It's so cool! It lets you send packets back and forth." No, I said: "If you get the Internet, you can call us and see your grandchildren on the screen. And if we have new pictures, you'll be able to see them immediately after we send them. And you can read about everything you can possibly imagine …"

Thus, the Internet does not create value through its existence alone. It creates value by enabling us to do the things we want to do, do things we never knew we wanted to do, or do things more efficiently. Applications are the tools that let us realize this value, in all areas of society. And by protecting the factors that have fostered application innovation in the past, we can make sure the Internet will be even more useful and valuable in the future.

Thank you for your attention.