Gmail by Google BETA

David Backus <david.backus@gmail.com>

notes on risk-adjusted distributions - and the disasters notebook

12 messages

Ian Martin <iwrmartin@gmail.com> **Tue, Jan 27, 2009 at 10:20 PM**
To: David Backus <dbackus@stern.nyu.edu>, Mikhail Chernov <mchernov@london.edu>

Hi guys,

Here, I hope, is something more substantial for you to get your teeth into! :-) Some notes on risk-neutral stuff in discrete time that, I think, suggest that we should be able to do everything in discrete time rather than referring people to an continuous-time appendix. (I hope my calculations are right - I have been looking at this algebra for a while now so I may have missed something.)

I'm also attaching that Mathematica notebook I promised you in my last email but which, I now see, I forgot to attach!

Cheers,
Ian

--
Ian Martin
Assistant Professor of Finance
Graduate School of Business
Stanford University

http://www.stanford.edu/~iwrm

3 attachments

disasters.nb
24K

notes on risk-neutral distributions.tex
9K

notes on risk-neutral distributions.pdf
84K

Mikhail Chernov <mchernov@london.edu> **Wed, Jan 28, 2009 at 1:24 AM**
To: Ian Martin <iwrmartin@gmail.com>, David Backus <dbackus@stern.nyu.edu>

Dear Ian,

This stuff can definitely be done in discrete time. The functional form

of power utility is tractable enough to allow for this. So this is all good. My personal challenge was to get similar expressions for the "finance" case. This is the case, where instead specifying SDF, one specifies p* directly. So, this is the stuff from the app. A.3.2. Once one allows for volatility of jumps to differ between the two probability measures, the tractability disappears in discrete time (or so it appears to me). Maybe the issue is the same: I was looking at the Bernoulli approximation to the Poisson process. What do you think?

Mike

[Quoted text hidden]

This email has been scanned by the MessageLabs Email Security System on behalf of the London Business School community.
For more information please visit http://www.messagelabs.com/email

---

Ian Martin <iwrmartin@gmail.com> **Wed, Jan 28, 2009 at 2:16 AM**
To: Mikhail Chernov <mchernov@london.edu>
Cc: David Backus <dbackus@stern.nyu.edu>

Hi Mike,

Yes, I don't think p* can be specified totally arbitrarily, eg jump volatility can't be allowed to differ between measures. But this isn't surprising in a sense: after all, even in the simple Brownian motion case (ignoring jumps) you can change the drift, but you can't change the volatility of the Brownian motion when you change measure----this is Girsanovs theorem, unless I am misunderstanding something?

Based on the discrete time calculations I sent you, it looks as though the volatility of both epsilon (the completely familiar normal shock) and the volatility of the disaster size can't change in the new measure, but that each of the other three parameters (mean of epsilon, arrival rate of Poisson process, and mean disaster size) can change arbitrarily.

Also, though this is somewhat off-topic, I still think there are important issues with how we present this stuff... I think most macro people will find the idea of just "writing down" a risk-neutral process for consumption growth extremely confusing unless we explain it really well.

Cheers,

Ian

[Quoted text hidden]

Mikhail Chernov <mchernov@london.edu> **Wed, Jan 28, 2009 at 2:47 AM**
To: Ian Martin <iwrmartin@gmail.com>
Cc: David Backus <dbackus@stern.nyu.edu>

No! That's the thing: generalization of Girsanov to jumps allows for the volatility of the jump size to change. One cannot change the volatility of the Brownian motion because it is observable in continuous-time. This is not the case for jumps. BCJ in JF gives a couple of references to math books/papers on this (p. 1458, first full paragraph).

Also, I think I now understand what you mean under incorrect computations in (19), etc. The security itself is weird: it starts to exist only when there is a jump. Otherwise, if we go down the route that you suggest, we will have jump intensity mixed with the jump mean. This is OK, but it would be harder to compare with power utility which delivers separate predictions for intensity and mean.

-----Original Message-----
From: Ian Martin [mailto:iwrmartin@gmail.com]

[Quoted text hidden]

Mikhail Chernov <mchernov@london.edu> **Wed, Jan 28, 2009 at 4:41 AM**
To: iwrmartin@gmail.com
Cc: dbackus@stern.nyu.edu

Regarding the risk-neutral motivation...

Can't we motivate it by the usual argument: in the risk-neutral world expected return on ALL securities is equal to the risk-free rate. This requirement locks us into a certain specification of the drift of the index returns.

Or do you want to motivate a different aspect of this thing?

Sent using BlackBerry

[Quoted text hidden]

[Quoted text hidden]

David Backus <dbackus@stern.nyu.edu> **Wed, Jan 28, 2009 at 8:44 AM**
To: Ian Martin <iwrmartin@gmail.com>
Cc: Mikhail Chernov <mchernov@london.edu>

This is great. I'll get back to work on the paper, starting at the beginning and working up to this. More later.

[Quoted text hidden]

--

http://whitepapers.stern.nyu.edu/home.html

Ian Martin <iwrmartin@gmail.com> **Wed, Jan 28, 2009 at 8:02 PM**

To: Mikhail Chernov <mchernov@london.edu>
Cc: David Backus <dbackus@stern.nyu.edu>

(I'm going to reply to the emails one at a time...)

On Tue, Jan 27, 2009 at 11:47 PM, Mikhail Chernov <mchernov@london.edu> wrote:
> No! That's the thing: generalization of Girsanov to jumps allows for the
> volatility of the jump size to change. One cannot change the volatility
> of the Brownian motion because it is observable in continuous-time. This
> is not the case for jumps. BCJ in JF gives a couple of references to
> math books/papers on this (p. 1458, first full paragraph).

Interesting. And, now that you tell me, it seems intuitive! :-) In fact, then, this may be something we can sell as an *advantage* of imposing some (power utility) structure on the model: it gets us some extra identification. (I have never been into econometrics so I hope I am using this terminology correctly....) In particular, we can beat ourselves up for imposing power utility but then console readers with the fact that it lets us identify the jump variance exactly. In terms of the structure-versus-generality tradeoff, this may be a tradeoff we're willing to make. In particular, it may provide a sense in which the equilibrium approach of the paper starts to show us how to interpret/ restrict the results that come out of more black-box (to my eyes.... may just be because of my ignorance!) approaches to option pricing.

> Also, I think I now understand what you mean under incorrect
> computations in (19), etc. The security itself is weird: it starts to
> exist only when there is a jump. Otherwise, if we go down the route that
> you suggest, we will have jump intensity mixed with the jump mean. This
> is OK, but it would be harder to compare with power utility which
> delivers separate predictions for intensity and mean.

Yes, exactly - but I think that it's still intuitive, and worth putting in to illustrate the point of how we identify fundamental parameters from asset prices, even if we have to point out that the price can be decomposed into an intensity piece and a mean size piece.

[Quoted text hidden]

--

[Quoted text hidden]

Ian Martin <iwrmartin@gmail.com> **Wed, Jan 28, 2009 at 8:07 PM**

To: Mikhail Chernov <mchernov@london.edu>
Cc: dbackus@stern.nyu.edu

Actually all I meant is that the risk-neutral logic is just very confusing for people who don't encounter it all the time - unless it's explicitly occurring as part of replicating-portfolio-type logic. But I think it would be fine if, for example, we lead in by reminding people of the idea of state prices, p* = p x m normalized by the riskless rate, and then derive p* explicitly given our assumptions on m. I think where people get confused is if you just write down: here's p ... , here's p* .... : then there's the question about what

you're building into your assumptions about p*, in other words what it means in economic terms. I also find the material on pricing simple "jump securities" etc is very helpful from this point of view.

[Quoted text hidden]

--

[Quoted text hidden]

Mikhail Chernov <mchernov@london.edu> **Thu, Jan 29, 2009 at 1:42 AM**
To: iwrmartin@gmail.com
Cc: dbackus@stern.nyu.edu

Ok!

Sent using BlackBerry

-----Original Message-----
From: Ian Martin <iwrmartin@gmail.com>
To: Mikhail Chernov

[Quoted text hidden]

[Quoted text hidden]

Mikhail Chernov <mchernov@london.edu> **Thu, Jan 29, 2009 at 1:54 AM**
To: Ian Martin <iwrmartin@gmail.com>
Cc: David Backus <dbackus@stern.nyu.edu>

-----Original Message-----
From: Ian Martin [mailto:iwrmartin@gmail.com]
Sent: Thursday, January 29, 2009 1:02 AM
To: Mikhail Chernov
Cc: David Backus
Subject: Re: notes on risk-adjusted distributions - and the disasters notebook

[Quoted text hidden]

Sorry. This ain't gonna work. Having different jump volatilities is crucial to generating the smile. The only other way to get the smile is to have a model with stochastic volatility and jumps in stochastic volatility. This is a much better model, but it will kill Barro right away and this is why we are not discussing it. I am sorry for self indulgence, but see figure 5 in BCJ-JF and the discussion around it. Also see figure 3 in the January ALT: a change in delta generates the curvature required to match the shape of a smile.

> Also, I think I now understand what you mean under incorrect

> computations in (19), etc. The security itself is weird: it starts to

> exist only when there is a jump. Otherwise, if we go down the route that

> you suggest, we will have jump intensity mixed with the jump mean. This

> is OK, but it would be harder to compare with power utility which

> delivers separate predictions for intensity and mean.

Yes, exactly - but I think that it's still intuitive, and worth

putting in to illustrate the point of how we identify fundamental

parameters from asset prices, even if we have to point out that the

price can be decomposed into an intensity piece and a mean size piece.

This is fine.

[Quoted text hidden]

Ian Martin <iwrmartin@gmail.com> **Thu, Jan 29, 2009 at 4:07 AM**
To: Mikhail Chernov <mchernov@london.edu>
Cc: David Backus <dbackus@stern.nyu.edu>

That's interesting. I was just looking at BCJ. It seems to me, though, that in your fully unconstrained model SVCJ, you don't get much improvement by allowing for a different jump size variance (so if you were penalizing the addition of extra free parameters, it might be desirable not to allow for a free jump size variance.... of course this is hard to formalize without a proper definition of "penalizing"), so perhaps it is merely compensating statistically for some other dimension that we're ignoring (just as we're ignoring the fact that volatility is time-varying).

More generally, my prior is that the high prices of out-of-the-money calls have nothing to do with disasters. Presumably there's something else going on there - but is that the topic of the paper?

[Quoted text hidden]

--

[Quoted text hidden]

Mikhail Chernov <mchernov@london.edu> **Thu, Jan 29, 2009 at 5:03 AM**
To: Ian Martin <iwrmartin@gmail.com>
Cc: David Backus <dbackus@stern.nyu.edu>

Sure, SVCJ beats SVJ hands down and one does not need changing volatility of the jump in prices. However,
1. In the SVCJ model one needs the mean jump in the variance to be different across the two prob. measures. Because of the exponential variance jump specification, the same parameter controls the volatility of the jump size, i.e. one ends up going into 2nd moments anyways.
2. It was a conscious choice to go with an iid model because Barro's stuff is iid. Otherwise, there is nothing to talk about: financial markets tell us that volatility should be stochastic and with jumps. This rejects his model. Done.

Regarding your second point, thinking about the magnitude of volatility is important because it affects both upside and downside. If it is too small, the downside, i.e. the disaster risk, is not strong.

-----Original Message-----
From: Ian Martin [mailto:iwrmartin@gmail.com]

[Quoted text hidden]