Gibbs estimation of microstructure models: Teaching notes

Joel Hasbrouck

```
Text @ Style["Evaluated " <> DateString["DateTime"], "Subsubtitle"]
```

Evaluated Tuesday 30 March 2010 11:51:28

1. Introduction

This note discusses Gibbs estimation of the Roll model and various modifications. The goal is a more discursive and heuristic treatment of material covered in Hasbrouck (2009). Other applications of Gibbs samplers in market microstructure include Hasbrouck (1999) and Ball and Chordia (2001).

The techniques discussed here follow an approach that relies on simulation to characterize model parameters. Applied to microstructure models, there are three key elements:

- *Bayesian analysis*
- *Simulation*
- *Characterization of microstructure data generating processes by their conditional probabilities.*

Specifically:

Bayesian analysis

The models are estimated in a Bayesian framework. The differences between Bayesian and classical analysis are continually debated and discussed. The framework here is definitely Bayesian, but it should be noted that even if one doesn't buy the full Bayesian philosophy, the techniques discussed here can be motivated on grounds of estimation simplicity and computational efficiency.
This is an unusual claim. Bayesian analyses are usually more complex (both conceptually and computationally) than their classical counterparts. This is sometimes cited by Bayesian adherents as the prime barrier to their widespread adoption. Most microstructure models, though, are dynamic (over time) and they have latent (hidden, unobservable) quantities. The classic Roll model is dynamic, and the trade direction indicator ("buy or sell") variables are not observed.

Dynamic latent variable models can be formulated in state-space form and estimated via maximum likelihood. The latent variables are often non-Gaussian (e.g., again, the trade direction indicator variables), and if one wants to go beyond the techniques of multivariate linear models (like VARs), estimation involves nonlinear filtering. The Gibbs estimates are usually quicker and simpler.

There are presently a number of Bayesian statistics textbooks available. In my opinion the most useful for financial econometricians, are those that discuss econometrics from a Bayesian perspective. Lancaster (2004) and Geweke (2005) are both excellent. Lancaster's treatment is particularly accessible; Geweke presents more results. Nelson and Kim (2000) is a good introduction to the techniques in the context of a specific problem (regime switching models). In financial econometrics, the heaviest use of Bayesian simulation has been in modeling stochastic volatility. Shephard (2005) is a good survery of this area. Tanner (1996) and Carlin and Louis (2004) consider a broader range of Bayesian statistical tools and applications.

Simulation.

The output of a classical procedure (e.g., OLS) is usually a statement about the distribution of a parameter. E.g., "θ is asympototically normally distributed with mean $\overline{\theta}$ and variance σ_{θ}^2," where the mean and variance quantities are computed directly. But we could also characterize a distribution by a sample of draws from that distribution. This is what most of modern Bayesian analysis does. The output of the estimation procedures discussed here is a stream of random draws of the parameters of interest (conditional on the model and the data). From this stream we can construct an estimate of the full distribution (via kernel smoothing) or simply a summary measure (like the mean or median).

Among other things, simulation facilitates characterization of distributions for functions of random variables. For example, suppose that $x \sim N\left(\mu, \sigma^2\right)$ and we'd like to characterize the distribution of $y = f\,(x)$ where f is sufficiently complicated that we can't get closed-form results. We simply generate random values x_i and empirically examine the distribution of $y_i = f\,(x_i)$.

The link between simulation and Bayesian analysis is strong for the following reason. The distributions that arise in Bayesian analysis often describe many random variables (i.e., they are of high dimension). It also often happens that they have no closed form representation. Instead, they are characterized by simulation. The Gibbs procedure belongs to a class of random number generators called Markov Chain Monte Carlo (MCMC) techniques. They work by setting up rules for moving from one realization (draw) of the random variables to a subsequent realization. These draws are viewed as "states" in a Markov chain, and the rules define the transition probabilities. The limiting distribution of the states is identical to the distribution of the variables of interest, and is approximated by repeated application of the transition function.

Conditional probabilities

To set up a Gibbs estimate, we need to compute conditional densities for all of the unknowns (parameters *and* latent data). The conditional distributions for the parameters are usually identical to those found in many other applications (e.g., the normal Bayesian linear regression model). This note merely summarizes these distributions, refering the reader elsewhere for a fuller treatment. The conditional distributions for the latent data, though, are specific to the problem at hand. Although not particularly complicated, they are non-standard, and this note covers them in greater detail.

Programs

This note is written in *Mathematica*. Some *Mathematica* code and results are embedded. Most of the results, though, are computed using SAS code that is available on my web site. These programs are available in my ftp directory at http://pages.stern-.nyu.edu/~jhasbrou/Teaching/2010%20 PhD %20 Microstructure/PhDMicroSp2010/.

The programs make heavy use of SAS/IML ("Interactive Matrix Language"). This is not the language I've used for most of my papers, but it is widely available. Anyone who has a copy of SAS should be able to run the programs. These programs are not "industrial strength". I've played around with them in generating the results for this note, but they haven't been tested against all the things that might come up in, say, the CRSP daily file. I haven't done any performance benchmarks, but I suspect that they run slower than comparable code in OX or Matlab.

The main programs used here are:

RollGibbs2Trade case.sas
RollGibbsAnalyzeq.sas
RollGibbs01.sas
RollGibbsBeta01.sas

These programs call two macros: *RollGibbs.sas* and *RollGibbsBeta.sas.* These macros, in turn, make use of IML subroutines contained in a library called "imlstor". To set up this library, run the program *RollGibbsLibrary01.sas* (which contains the code for the subroutines).

2. Overview

This note illustrates the estimation approach for the Roll (1984) model of transaction prices. In this model, the "efficient price" (m_t) is assumed to follow a Gaussian random walk:

$$m_t = m_{t-1} + u_t \text{ where } u_t \sim N\left(0, \sigma_u^2\right)$$

The transaction price (p_t) is the efficient price, plus or minus a cost component that depends on whether the customer is buying or selling:

$$p_t = m_t + c\, q_t$$

where c is the cost parameter and $q_t = \pm 1$. (If the customer is buying, $q_t = +1$; if selling, $q_t = -1$). The trade prices are observed. The q_t and the m_t are not. Taking first differences:

$$\Delta p_t = c\, \Delta \mathrm{q}_t + u_t$$

This specification is important because if the $\Delta \mathrm{q}_t$ were known, this would be a simple regression.

Bayesian estimation of normal linear regressions is well understood. The discussion (in the next section) starts with a review of these procedures.

There are two parameters in this "regression": c (the coefficient) and σ_u^2. It is fairly easy to compute (in closed form) the posterior distributions $f\left(c \,|\, \sigma_u^2,\ p_1,\ ...,\ p_T\right)$ and $f\left(\sigma_u^2 \,|\, c,\ p_1,\ ...,\ p_T\right)$. It is not possible to compute in closed form the joint posterior $f\left(c, \sigma_u^2 \,|\ p_1,\ ...,\ p_T\right)$. This motivates the next section, which summarizes the Gibbs sampler.

The Gibbs procedure is illustrated by applying it to a special case of the Roll model, one in which c and σ_u^2 are known, but the q_t are not. The note then turns to a full estimation of the Roll model, and extensions.

3. *Mathematica* initializations

```
SetDirectory[NotebookDirectory[]];
```

```
<< MVN.m
```

```
<< Notation`
```

The following commands define symbolizations that are convenient for labeling things.

```
Symbolize[Anything_Rule]; Symbolize[Anything_Rules]
```

```
Off[General::"spell", General::"spell1"];
```

4. Bayesian analysis of the normal linear regression model

The basic Bayesian approach

Bayesian analysis consists of using a model and data to update prior beliefs.The revised beliefs are usually called posterior beliefs, or simply "the posterior". Let y denote the observed data, and let the model be specified up to a parameter θ (possibly a vector). Bayes theorem says :

$$f(\theta \mid y) = \frac{f(\theta, y)}{f(y)} = \frac{f(y|\theta)\, f(\theta)}{f(y)} \propto f(y \mid \theta)\, f(\theta)$$

- $f(\theta)$ is an assumed prior distribution for the parameter.
- $f(y \mid \theta)$ is the likelihood function for the observations, given a particular value of θ.

The use of $\propto$ ("is proportional to") reflects the fact that it is usually not necessary to compute $f(y)$, at least not by computing the marginal $f(y) = \int f(\theta, y)\, d\theta$. Instead, $f(y)$ can treated as a normalization constant, set so that the posterior integrates to unity.

Often a distribution of interest, say $f(x)$, can be written as $f(x) = k\, g(x)$, where $g(x)$ is a parsimonious function of x and k is some scaling factor. k might in fact be very complicated, possibly depending on other random variables and implicitly incorporating other distribution functions, but for purposes of characterizing the distribution of x, it is constant. In this case, $g(x)$ is said to be the *kernel* of $f(x)$.

Bayesian estimation of the normal linear regression model

The normal regression model is:

$$\underset{N\times 1}{y} = \underset{N\times K}{X}\ \underset{K\times 1}{\beta} + \underset{N\times 1}{u} \text{ where } u \sim N(0, \Omega_u)$$

X is a matrix of covariates (explanatory variables) possibly including a constant; β is the coefficient vector.

Estimation of coefficients (given the error variance)

Assume for the moment that σ_u^2 is known. It is particularly convenient to assume a multivariate normal prior distribution for the coefficients:

$$\beta \sim N\left(\mu_\beta^{\text{Prior}}, \Omega_\beta^{\text{Prior}}\right)$$

The posterior distribution, $f(\beta \mid y)$is

$$N\left(\mu_\beta^{\text{Post}}, \Omega_\beta^{\text{Post}}\right)$$

where $\mu_\beta^{\text{Post}} = D\, d$; $\quad \Omega_\beta^{\text{Post}} = D^{-1}$; $\quad D^{-1} = X'\, \Omega_u^{-1}\, X + \left(\Omega_\beta^{\text{Prior}}\right)^{-1}$; $\quad d = X'\, \Omega_u^{-1}\, y + \left(\Omega_\beta^{\text{Prior}}\right)^{-1} \mu_\beta^{\text{Prior}}$

As $\Omega_\beta^{\text{Prior}}$ increases in magnitude, the posterior mean and variance converge toward the usual OLS values.
In this case, both the prior and posterior have the same form (multivariate normal). Such a prior is said to be conjugate.

Simulating the coefficients

We'll often have to make a random draw from the coefficient distribution. To make a random draw from $\underset{n\times 1}{x} \sim \text{MVN}(\mu, \Omega)$:

- Compute the Cholesky factorization F : $\Omega = F'\, F$, where F is an upper triangular matrix.

Draw $\underset{n\times 1}{z}$ where the z_i are i.i.d. $N(0, 1)$.

- Set the random draw as $x = \mu + F' z$

Restrictions on the prior

The economic model sometimes imposes bounds on the coefficients. For example, in the Roll model, we'll often want to require $c > 0$. Suppose that the coefficient prior is

$$\beta \sim N\left(\mu_\beta^{\text{Prior}}, \Omega_\beta^{\text{Prior}}\right),\ \underline{\beta} < \beta < \overline{\beta}$$

Note that when we write this, μ_β^{Prior}, $\Omega_\beta^{\text{Prior}}$ denote the formal parameters of the normal density. But since the distribution is truncated, they no longer denote the mean and covariance of the density.

With this prior, the posterior is simply $N\left(\mu_\beta^{\text{Post}}, \Omega_\beta^{\text{Post}}\right)$, with μ_β^{Post} and $\Omega_\beta^{\text{Post}}$computed as described above, restricted to the space $\underline{\beta} < \beta < \overline{\beta}$.

Simulation from restricted normals.

First suppose that we want to make a random draw z from a standard normal density, restricted to the interval $\underline{z} < z < \overline{z}$. The procedure is:

- Compute $\underline{\pi} = \Phi(\underline{z})$ and $\overline{\pi} = \Phi(\overline{z})$, where Φ is the c.d.f. of the standard normal.
- Draw u from the uniform distribution over $(\underline{\pi}, \overline{\pi})$
- Set $z = \Phi^{-1}(u)$.

Now suppose that we want to make a bivariate random draw from $x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \sim N(\mu, \Omega)$, $\underline{x} < x < \overline{x}$.

- Compute the Cholesky factorization $F:\ \Omega = F' F$, where F is an upper triangular matrix.
- Set $\underline{z} = \frac{\underline{x_1} - \mu_1}{F_{11}}$ and $\overline{z} = \frac{\overline{x_1} - \mu_1}{F_{11}}$
- Draw z_1 from the standard normal density, restricted to $(\underline{z}, \overline{z})$. Then $x_1 = \mu_1 + F_{11} z_1$ will have the properties required of x_1.
- Set $\eta = F_{11} z_1$.
- Set $\underline{z} = \frac{\underline{x_2} - \mu_2 - \eta}{F_{22}}$ and $\overline{z} = \frac{\overline{x_2} - \mu_2 - \eta}{F_{22}}$
- Draw z_2 from the standard normal density, restricted to $(\underline{z}, \overline{z})$. Then $x_2 = \mu_2 + F_{22} z_2$ will********** have the properties required of x_2.
- The random draw as $x = \mu + F' z$ will have the required joint properties

This method may be generalized to higher dimensions. (See Hajivassiliou, V., D. McFadden and P. Ruud (1996). Simulation of multivariate normal rectangle probabilities and their derivatives - Theoretical and computational results. Journal of Econometrics 72(1-2): 85-134)

Estimation of error variance (given the coefficients)

Assuming that β is known, it is convenient to specify a inverted gamma prior for σ_u^2. One way of writing this is:

- $\frac{1}{\sigma_u^2} \sim \Gamma\left[a^{\text{Prior}}, b^{\text{Prior}}\right]$

Then the posterior is

- $\frac{1}{\sigma_u^2} \,\Big|\, y \sim \Gamma\left[a^{\text{Post}}, b^{\text{Post}}\right]$

where $a^{\text{Post}} = a^{\text{Prior}} + \frac{N}{2}$ and $b^{\text{Post}} = b^{\text{Prior}} + \frac{\sum_{i=1}^{N} u_i^2}{2}$

The u_i are the regression residuals $u = y - X\,\beta$.

Further notes

The density of a gamma variate x with parameters a and b is:

```
PDF[GammaDistribution[a, λ], x] /. λ → 1 / b
```

$$\frac{\left(\frac{1}{b}\right)^{-a}\, e^{-b\,x}\, x^{-1+a}}{\text{Gamma}[a]}$$

```
Mean[GammaDistribution[a, λ]] /. λ → 1 / b
```

$$\frac{a}{b}$$

Note: In the statistics literature, the Gamma distribution with parameters a and b is usually expressed as immediately above. *Mathematica* parameterizes the distribution with the second parameter expressed as an inverse.

```
PDF[GammaDistribution[a, λ], x] /. λ → 1 / b /. x → 1 / z
```

$$\frac{\left(\frac{1}{b}\right)^{-a}\, e^{-\frac{b}{z}}\, \left(\frac{1}{z}\right)^{-1+a}}{\text{Gamma}[a]}$$

5. The Gibbs recipe

The Gibbs procedure is a vehicle for simulating from a complicated joint distribution $f\,(x_1,\ \ldots,\ x_n)$, possibly one that possesses no closed form representation.

The draws are constructed by iterating over the *full conditional distributions*:

$$f\,(x_1 \mid x_2,\ \ldots,\ x_n)$$
$$f\,(x_2 \mid x_1, x_3,\ \ldots,\ x_n)$$
$$\ldots$$
$$f\,(x_n \mid x_1,\ \ldots,\ x_{n-1})$$

Let $x = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix}$

Each iteration of the Gibbs sampler is called a sweep.
Let $x^{[i]}$ denote the value of x at the conclusion of the ith sweep.
The procedure is:

Initialization. Set $x^{[0]}$ to any feasible value.

Sweep i:
Given $x^{[i-1]}$:

- Draw $x_1^{[i]}$ from $f\left(x_1 \mid x_2^{[i-1]},\ ...,\ x_n^{[i-1]}\right)$
- Draw $x_2^{[i]}$ from $f\left(x_2 \mid x_1^{[i]}, x_3^{[i-1]},\ ...,\ x_n^{[i-1]}\right)$
  ...
- Draw $x_n^{[i]}$ from $f\left(x_n \mid x_1^{[i]},\ ...,\ x_{n-1}^{[i]}\right)$

Repeat

In the limit, as $i \to \infty$, $x^{[i]}$is distributed as $f(x)$.

Notes

- The $x^{[i]}$ are *not* independently distributed: $x^{[i]}$ takes $x^{[i-1]}$ as its starting point. If the degree of dependence is high, a large number of sweeps may be needed to ensure proper mixing.
- The dependence may affect the calculation of some summary statistics. Think of the analogy to standard time series analysis. If $z_1,\ ...,\ z_T$ are a sample of stationary time series data, $\Sigma z_i / T$ is a consistent estimate of $E[z_t]$. The standard error of this estimate, however, must be corrected for dependence.
- Convergence may be an issue. It is useful to graph the full sequence of draws.
- In analyzing the sequence of draws, it is common to throw out a few initial draws, so as to reduce the dependence on starting values. These discarded draws are sometimes called *burn in* draws.
- The Gibbs sampler also works when multiple variables are drawn at once. We might, for example, draw $x_1^{[i]}$ and $x_2^{[i]}$ from $f\left(x_1, x_2 \mid x_3^{[i-1]},\ ...,\ x_n^{[i-1]}\right)$. This *block sampling* is often more computationally efficient.

Application to the normal regression model

From earlier results, we have $f\left(\beta \mid y, \sigma_u^2\right)$ and $f\left(\sigma_u^2 \mid \beta, y\right)$. To obtain the full posterior $f\left(\beta, \sigma_u^2 \mid y\right)$ via the Gibbs procedure:

Initialize $\sigma_u^{2[0]}$ to any feasible value. The `i`th sweep of the sampler is:

- Draw $\beta^{[i]}$ from $f\left(\beta \mid y, \sigma_u^{2[i-1]}\right)$. (This will be a draw from a multivariate normal posterior.)
- Draw $\sigma_u^{2[i]}$ from $f\left[\sigma_u^2 \mid y, \beta^{[i]}\right)$. (That is, draw $1 / \sigma_u^{2[i]}$ from the gamma posterior.)

Proceed, iteratively drawing $\sigma_u^{2[i]}$ and $\beta^{[i]}$.

Notes

The $f\left(\beta, \sigma_u^2 \mid y\right)$ is an exact small-sample distribution.

We now return to ...

6. The Roll model

Recall that we're working with the price change specification:

$$\Delta p_t = c\, \Delta q_t + u_t$$

The sample is $p_1,\ p_2,\ ...,\ p_T$, and there are $T - 1$ price changes.

The unknowns are the parameters c and σ_u^2, and the latent data: $q_1,\ ...,\ q_T$.
In the Bayesian perspective, parameters and latent data are treated identically, and "estimated" in a similar fashion.
We don't need to construct priors for the q_t. We can use the ones we assumed in the data generating process: $q_t = \pm 1$ with equal probabilities.
The prior on c is $c \sim N\left(\mu^{\text{Prior}},\ \Omega^{\text{Prior}}\right)$ restricted to $c > 0$. (I often take $\mu^{\text{Prior}} = 0$ and $\Omega^{\text{Prior}} = 1$. Remember that these are parameters of the truncated distribution, not the true mean and variance.)
The prior on σ_u^2 is inverted gamma with parameters a and b. (I often take $a = b = 10^{-6}$.)
The Gibbs sampler will look like this:

- Initialize $c^{[0]}$, $\sigma_u^{2[0]}$ and $q_1^{[0]},\ ...,\ q_T^{[0]}$ to any feasible values.
  (I usually take $q_1 = 1$; $q_t = \text{Sign}\left(\Delta p_t\right)$ if $\Delta p_t \neq 0$, $q_t = q_{t-1}$ if $\Delta p_t = 0$. For US equities, $c = 0.01$ ("1%") is a good ballpark figure, if we're working in logs, and $\sigma_u^2 = 0.01^2$).

For the ith sweep of the sampler:

- Estimate the price change specification as a regression, assuming that $q_t = q_t^{[i-1]}$ and that $\sigma_u^2 = \sigma_u^{2[i-1]}$.
  Construct the posterior for c, and draw $c^{[i]}$ from this posterior.
- Using $c^{[i]}$, compute the residuals from the regression. Construct the posterior for σ_u^2 and draw $\sigma_u^{2[i]}$ from this posterior.
- Draw $q_1^{[i]}$ from $f\left(q_1 \mid c^{[i]}, \sigma_u^{2[i]}, q_2^{[i-1]}, q_3^{[i-1]},\ ...,\ q_T^{[i-1]}\right)$
  Draw $q_2^{[i]}$ from $f\left(q_2 \mid c^{[i]}, \sigma_u^{2[i]}, q_1^{[i]}, q_3^{[i-1]},\ ...,\ q_T^{[i-1]}\right)$
  ...
  Draw $q_T^{[i]}$ from $f\left(q_T \mid c^{[i]}, \sigma_u^{2[i]}, q_1^{[i]}, q_2^{[i]},\ ...,\ q_{T-1}^{[i]}\right)$

The first two steps of the sweep are, as discussed, standard Bayesian procedures. We now turn to the third step: simulating the trade direction indicators

7. Estimating the q_t using a Gibbs sampler.

In this section, we'll be taking c and σ_u^2 as known. We'll first look at the simple case where $T = 2$. We can get closed-form results here, so we don't really need a Gibbs procedure, but it's a good starting point.

The distribution of q_1 and q_2 when T=2

```
Symbolize[p_]; Symbolize[q1]; Symbolize[q2]; Symbolize[σ_]; Symbolize[u2];
```

```
$TextStyle = {FontFamily → "Times", FontSize → 12, LineSpacing → {1, -1}};
```

Graph

```
s = 0.05; hc = 0; hl = 0.02;
```

Suppose:

```
p1 = 1; p2 = 1.8; c = .5;
```

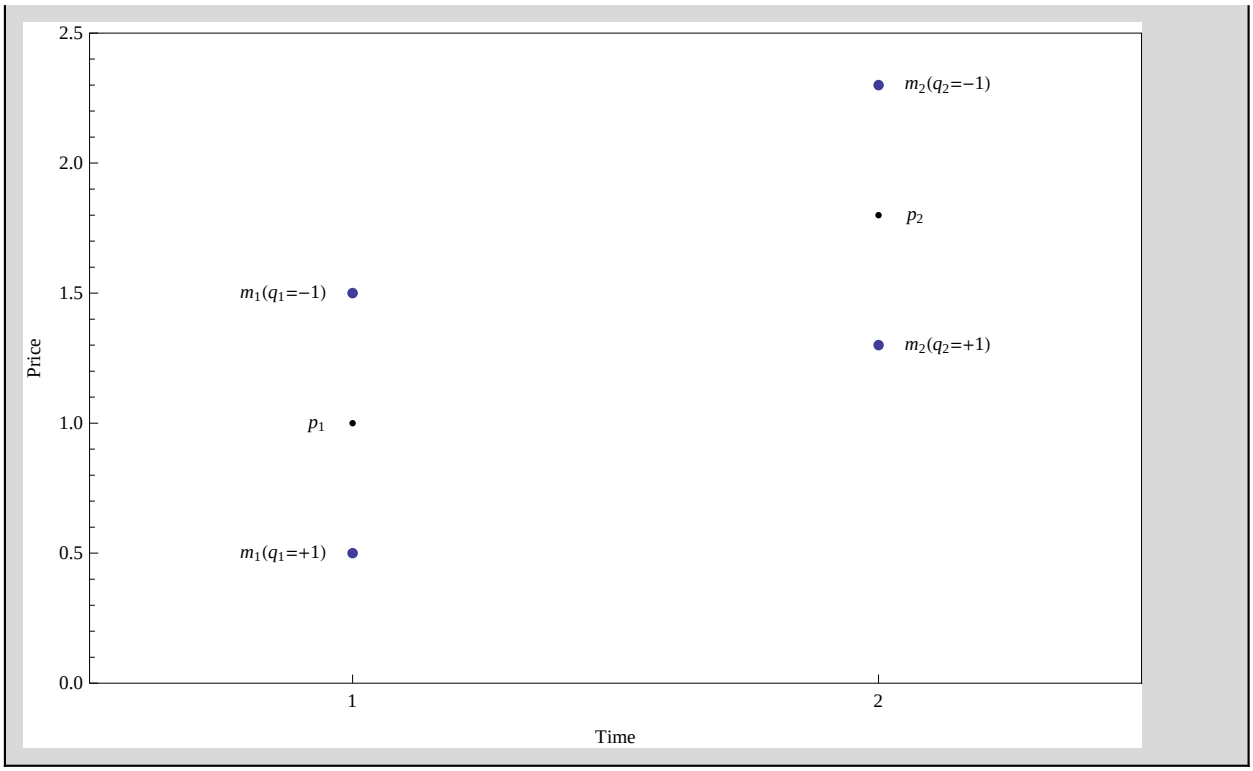

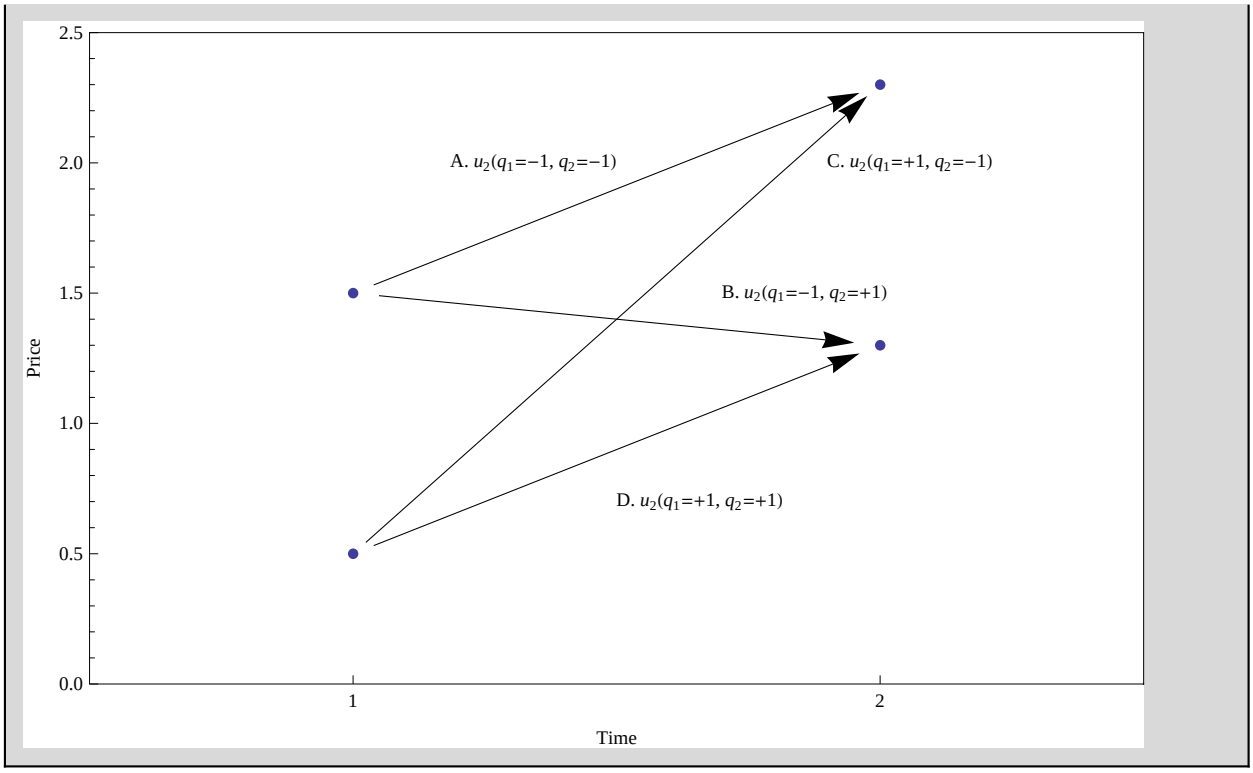

Given p_1 and c, a choice of q_1 determines $\mathtt{m}$, and similarly for p_2. Therefore, setting q_1 and q_2 fixes $u_2 = m_2 - m_1$.

Since $q_1, q_2 \in \{-1, +1\}$, there are four possible paths (values) for u_2. These are labeled A, B, C and D in the figure. Since low values of $|u_2|$ are more likely than high values, the relative lengths of these paths indicate the relative likelihood of the (q_1, q_2) realizations that determined them:

Intuitively, path B is the shortest, so it is most likely that $q_1 = -1,\ q_2 = +1$.

Path C is the longest, so it is least likely that $q_1 = +1,\ q_2 = -1$.

Paths A and D are of equal length, corresponding to the realizations $q_1 = q_2 = -1$ and $q_1 = q_2 = +1$.

We now turn to a more exact treatment.

```
Clear[s, p₁, p₂, c];
```

The joint density of q_1 and q_2

The density function for u is $f(u_2) =$

```
PDF[NormalDistribution[0, σu], u]
```

$$\frac{e^{-\frac{u^2}{2\sigma_u^2}}}{\sqrt{2\pi}\ \sigma_u}$$

By rearranging the price change specification, $u_2 =$

```
u₂[q₁ : _, q₂ : _] := Δp₂ - c (q₂ - q₁)
```

```
u₂[q₁, q₂]
```

$$-c\ (-q_1 + q_2) + \Delta p_2$$

The probability $\Pr[q_1, q_2] \propto f(u_2[q_1, q_2])$

```
Pr_Rule = Pr[q₁ : _, q₂ : _] :> PDF[NormalDistribution[0, σu], u₂[q₁, q₂]] // Simplify;
```

```
Pr[q₁, q₂] /. Pr_Rule
```

$$\frac{e^{-\frac{(-c\ (-q_1+q_2)+\Delta p_2)^2}{2\sigma_u^2}}}{\sqrt{2\pi}\ \sigma_u}$$

The possible outcomes are:

```
Outcomes = {{-1, -1}, {-1, +1}, {+1, -1}, {+1, +1}};
TableForm[Outcomes, TableHeadings → {None, {q1, q2}}]
```

```
q1   q2
-1   -1
-1   1
1    -1
1    1
```

We'll normalize by the sum of the probabilities:

```
PrSum = Plus @@ Apply[Pr, Outcomes, {1}]
```

```
Pr[-1, -1] + Pr[-1, 1] + Pr[1, -1] + Pr[1, 1]
```

```
PrNRule = PrN[q1 : _, q2 : _] :↦ Pr[q1, q2] / PrSum;
```

The normalized probability is:

```
PrN[q1, q2] /. PrNRule
```

$$\frac{\text{Pr}[q_1, q_2]}{\text{Pr}[-1, -1] + \text{Pr}[-1, 1] + \text{Pr}[1, -1] + \text{Pr}[1, 1]}$$

For demonstration purposes, here are some values:

```
nValues = {Δp2 → .8, σu → 1, c → 0.5};
```

With thes values, the normalized probabilities are:

```
TableForm[Transpose[
  Append[Transpose[Outcomes], Apply[PrN, Outcomes, {1}] /. PrNRule /. PrRule /. nValues]],
 TableHeadings → {None, {q1, q2, "Probability"}}]
```

```
q1   q2   Probability
-1   -1   0.276061
-1   1    0.372643
1    -1   0.0752353
1    1    0.276061
```

Gibbs sampler

We can draw the trade direction indicator variables jointly in this case. There's no need to use a Gibbs sampler. But for illustration purposes, let's build it.

The required conditional probabilities are $\Pr[q_1 \mid q_2]$ and $\Pr[q_2 \mid q_1]$.

These may be computed directly from joint distribution given above, but it is usually computationally easier to work with the odds ratio.

For example, $u_2[q_1, q_2] =$

```
u₂[q₁, q₂]
```

```
-c (-q₁ + q₂) + Δp₂
```

Its density is:

```
f = PDF[NormalDistribution[0, σᵤ], u₂[q₁, q₂]]
```

$$\frac{e^{-\frac{(-c(-q_1+q_2)+\Delta p_2)^2}{2\sigma_u^2}}}{\sqrt{2\pi}\ \sigma_u}$$

So given q_2, the odds in favor of a buy at time 1 are $\text{Odds (Buy)} = \frac{\Pr[q_1=+1|...]}{\Pr[q_1=-1|...]} =$

```
OddsBuy1 = f /. q₁ → +1 // Simplify
```

$$e^{\frac{2c(cq_2-\Delta p_2)}{\sigma_u^2}}$$

Then $\Pr[\text{Buy}] = \frac{\text{Odds (Buy)}}{1+\text{Odds (Buy)}}$. We compute this probability and make a draw for q_1.

For the particular numeric values we worked with above $\left(\Delta p_2 = 0.8,\ c = 0.5,\ \sigma_u = 1\right)$, these odds, for $q_2 = +1$ and $q_2 = -1$ are:

```
OddsBuy1 /. Δp₂ → 0.8 /. c → 0.5 /. σᵤ → 1 /. q₂ → {+1, -1}
```

```
{0.740818, 0.272532}
```

So, for example, if $q_2 = +1$, $\Pr[q_1 = +1 \mid ...] =$

```
%[[1]] / (1 + %[[1]])
```

```
0.425557
```

Similarly, given q_1, the odds in favor of a buy at time 2 are $\frac{\Pr[q_2=+1|...]}{\Pr[q_2=-1|...]} =$

```
OddsBuy2 = f /. q2 → +1 / f /. q2 → -1 // Simplify
```

$$e^{\frac{2c\,(c\,q_1+\Delta p_2)}{\sigma_u^2}}$$

For the numeric values, these are (for $q_1 = +1$ and $q_1 = -1$):

```
OddsBuy2 /. Δp2 → 0.8 /. c → 0.5 /. σu → 1 /. q1 → {+1, -1}
```

```
{3.6693, 1.34986}
```

The Gibbs sampler involves the following steps:

We construct a series of realizations in the following fashion:

- Initialize $q_1^{[0]}$, $q_2^{[0]} = \pm 1$ (it doesn't matter which).
  The ith sweep involves the following steps:
- Draw $q_1^{[i]}$ from $\Pr\left[q_1 \mid q_2^{[i-1]}\right]$
- Draw $q_2^{[i]}$ from $\Pr\left[q_2 \mid q_1^{[i]}\right]$

After N sweeps we'll have a series of N simulated realizations: $q^{[0]}$, $q^{[1]}$, ..., $q^{[N]}$ where $q^{[i]} = \left(q_1^{[i]}, q_2^{[i]}\right)$.
In the limit, as $N \to \infty$, the distribution of $q^{[N]}$ is the joint distribution $\Pr[q_1, q_2]$

The sampler was implemented in a SAS program (*RollGibbs2TradeCase.sas*), which was run for 10,000 sweeps. The tabulated frequencies of the simulated draws were:

	N	PctN
outcome		
q1= 1 q2= 1	2,752	27.5
q1= 1 q2=-1	749	7.5
q1=-1 q2= 1	3,729	37.3
q1=-1 q2=-1	2,770	27.7
All	10,000	100.0

Compare these to the computed probabilities above.

Why do we ever need to use Gibbs sampler when we can compute path probabilities directly? We need to compute path probabilities for the entire sample. With two price observations, there are $2^2 = 4$ buy/sell paths. A year contains about 250 trading days. The number of buy/sell paths is:

```
2^250
```

```
1 809 251 394 333 065 553 493 296 640 760 748 560 207 343 510 400 633 813 116 524 750 123 642 650 624
```

The T-trade case

A sweep of the sampler will involve:

- Draw $q_1^{[i]}$ from $\Pr\left(q_1 \mid q_2^{[i-1]}, q_3^{[i-1]}, \ ..., \ q_T^{[i-1]}\right)$
- Draw $q_2^{[i]}$ from $\Pr\left(q_2 \mid q_1^{[i]}, q_3^{[i-1]}, \ ..., \ q_T^{[i-1]}\right)$

  ...
- Draw $q_T^{[i]}$ from $\Pr\left(q_T \mid q_1^{[i]}, q_2^{[i]}, \ ..., \ q_{T-1}^{[i]}\right)$

In general, $\Pr[q_t \mid \ ...]$depends only on the adjacent trades - those at times $t-1$ and $t+1$. So the sampler becomes:

- Draw $q_1^{[i]}$ from $\Pr\left(q_1 \mid q_2^{[i-1]}\right)$
- Draw $q_2^{[i]}$ from $\Pr\left(q_2 \mid q_1^{[i]}, q_3^{[i-1]}\right)$

  ...
- Draw $q_T^{[i]}$ from $\Pr\left(q_T \mid \ q_{T-1}^{[i]}\right)$

The first draw, for q_1, is the same as the draw for q_1 in the $T = 2$ case.

The last draw, for q_T, is the same as the draw for q_2 in the $T = 2$ case.

We now turn to the middle draws.

```
u_Rule = u_t_ :> Δp_t + c (q_t - q_t-1);
```

```
u_t+1 /. u_Rule
```

```
c (-q_t + q_1+t) + Δp_1+t
```

Since the u_t are assumed to be independent, the joint density $f\left(u_t, u_{t+1}\right) \propto$

```
PDF[NormalDistribution[0, σ_u], u_t] PDF[NormalDistribution[0, σ_u], u_t+1]
```

$$\frac{e^{-\frac{u_t^2}{2\sigma_u^2}-\frac{u_{1+t}^2}{2\sigma_u^2}}}{2\pi\sigma_u^2}$$

```
f = % /. u_Rule // Simplify
```

$$\frac{e^{-\frac{(-c\,q_{-1+t}+c\,q_t+\Delta p_t)^2+(-c\,q_t+c\,q_{1+t}+\Delta p_{1+t})^2}{2\sigma_u^2}}}{2\pi\sigma_u^2}$$

The odds ratio is:

$$\text{Odds} = \frac{\Pr[q_t{=}{+}1|...]}{\Pr[q_t{=}{-}1|...]}$$

```
Odds = f /. q_t → +1 / f /. q_t → -1 // Simplify
```

$$e^{\frac{2c\left(c\,q_{-1+t}+c\,q_{1+t}-\Delta p_t+\Delta p_{1+t}\right)}{\sigma_u^2}}$$

from which we may compute

$$\Pr[q_t = +1 \mid \ldots] = \frac{\text{Odds}}{1+\text{Odds}}$$

and make the desired draw.

Examples (SAS program *RollGibbs Analyze q.sas*)

I simulated twenty trades for a price process with σ_u = `0.01` and `c` = `0.01`, and then ran the Gibbs sampler for 2,000 sweeps to estimate the trade directions. Here is a plot of the transaction prices:

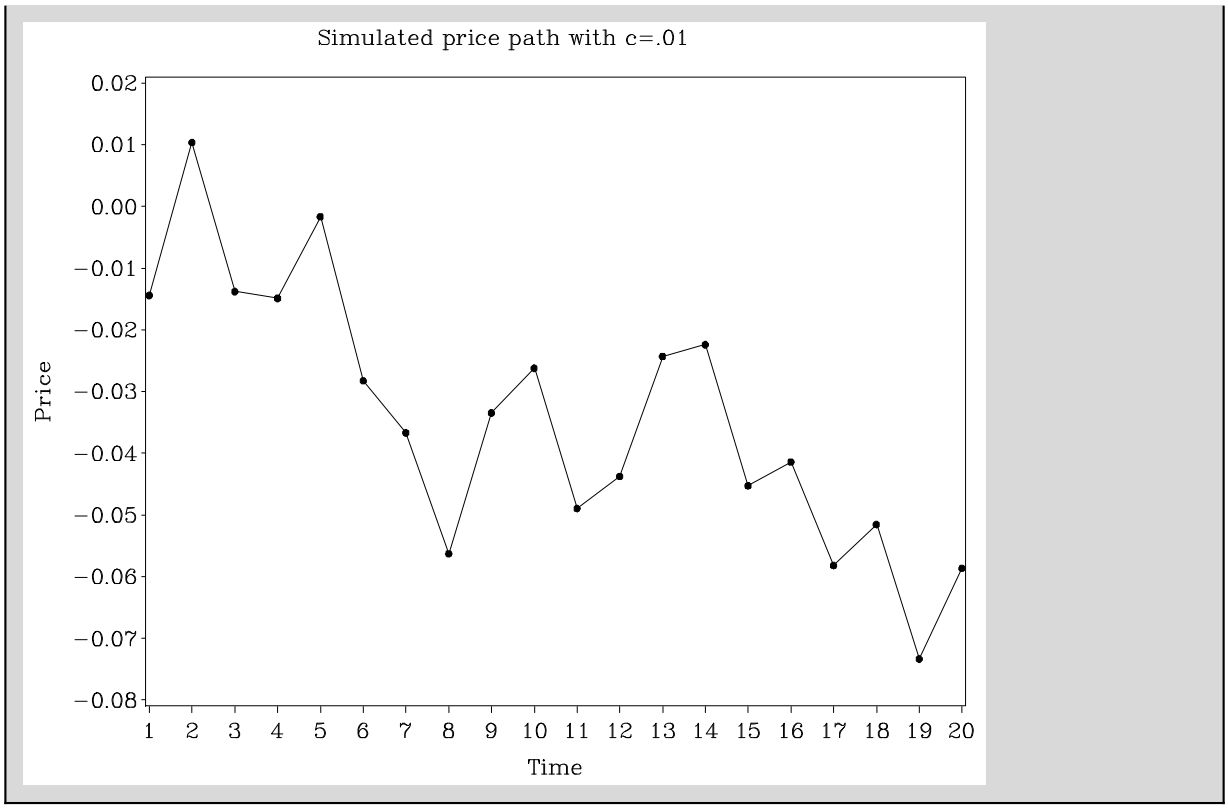

Below are the actual and average simulated trade directions. Actuals are indicated by a dot; estimated are indicated by the bars.

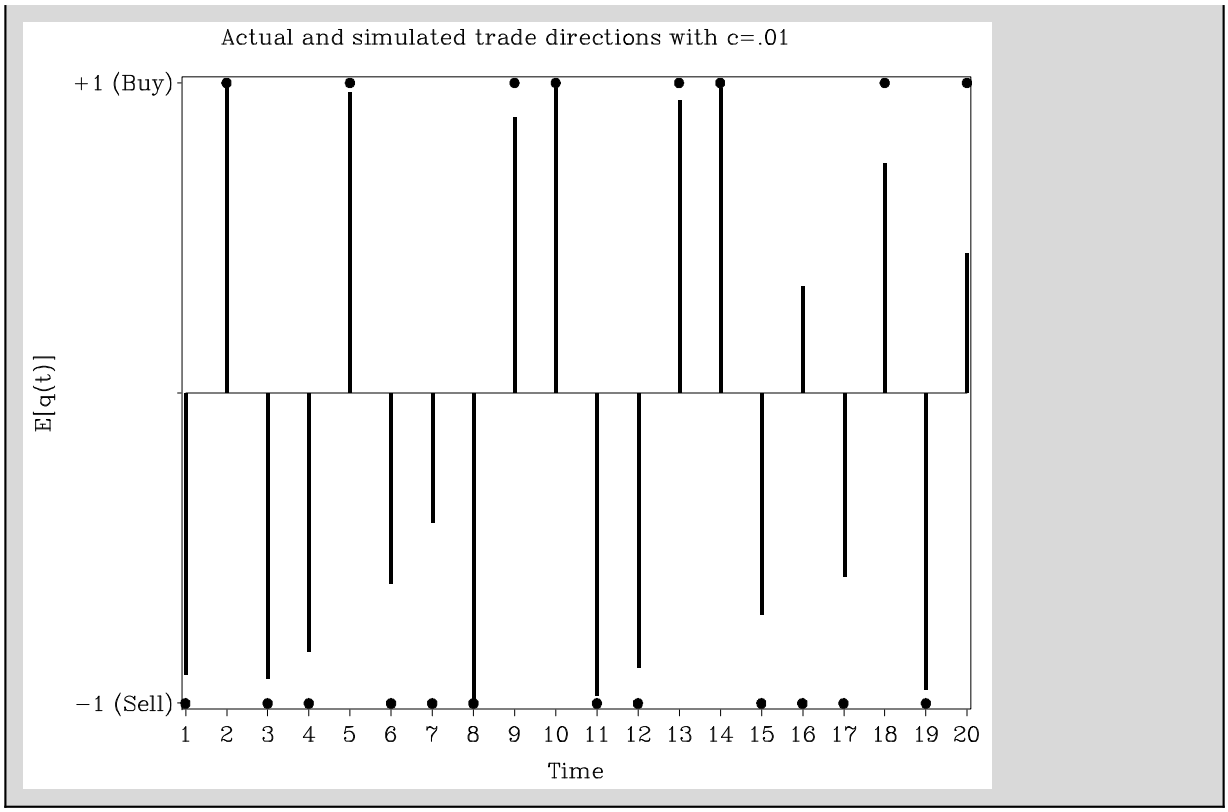

Note that estimates by and large agree with with actuals, at least in direction (q_{16} is the sole exception).

Now, consider the same analysis, with the cost parameter changed to `c = 0.005`, i.e., one-half the previous value. The figure below shows the prices. Notice that the bid-ask bounce is much less visually evident.

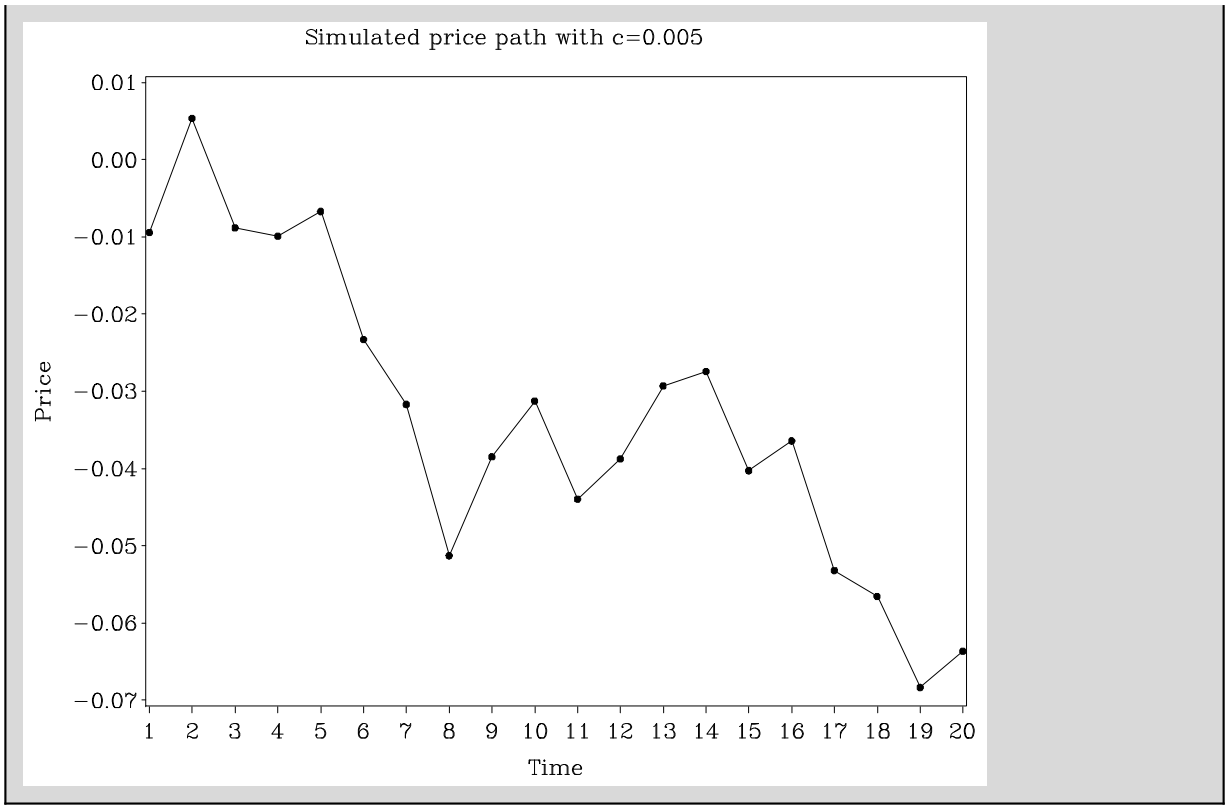

Here are the actual and estimated trade direction indicators:

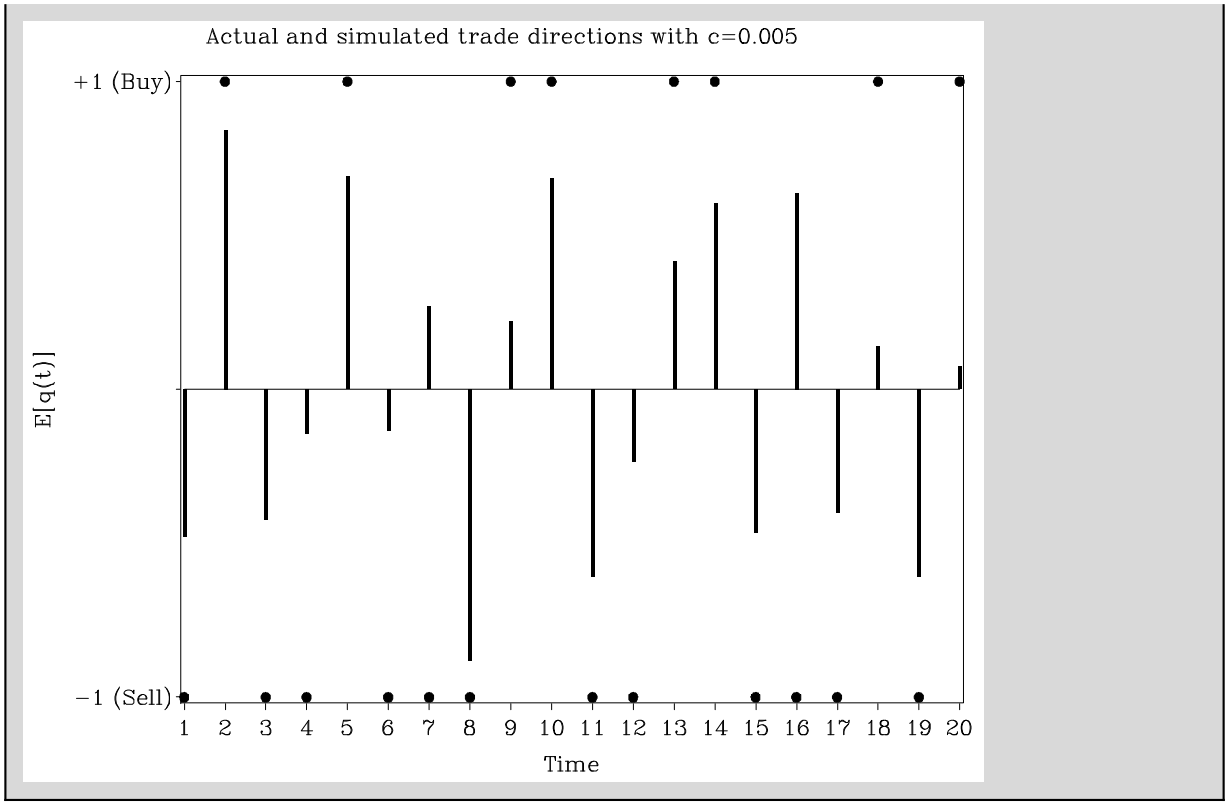

The buy/sell classification accuracy here is weaker. There are more directionally incorrect inferences, and the directions that are correctly identified are weaker. Just as we'd have a harder time picking out the buys and sells visually, the sampler has a tougher time classifying trades.

We'll now try things with σ_u = `0.01` and `c` = `0.001`. Here's the price path:

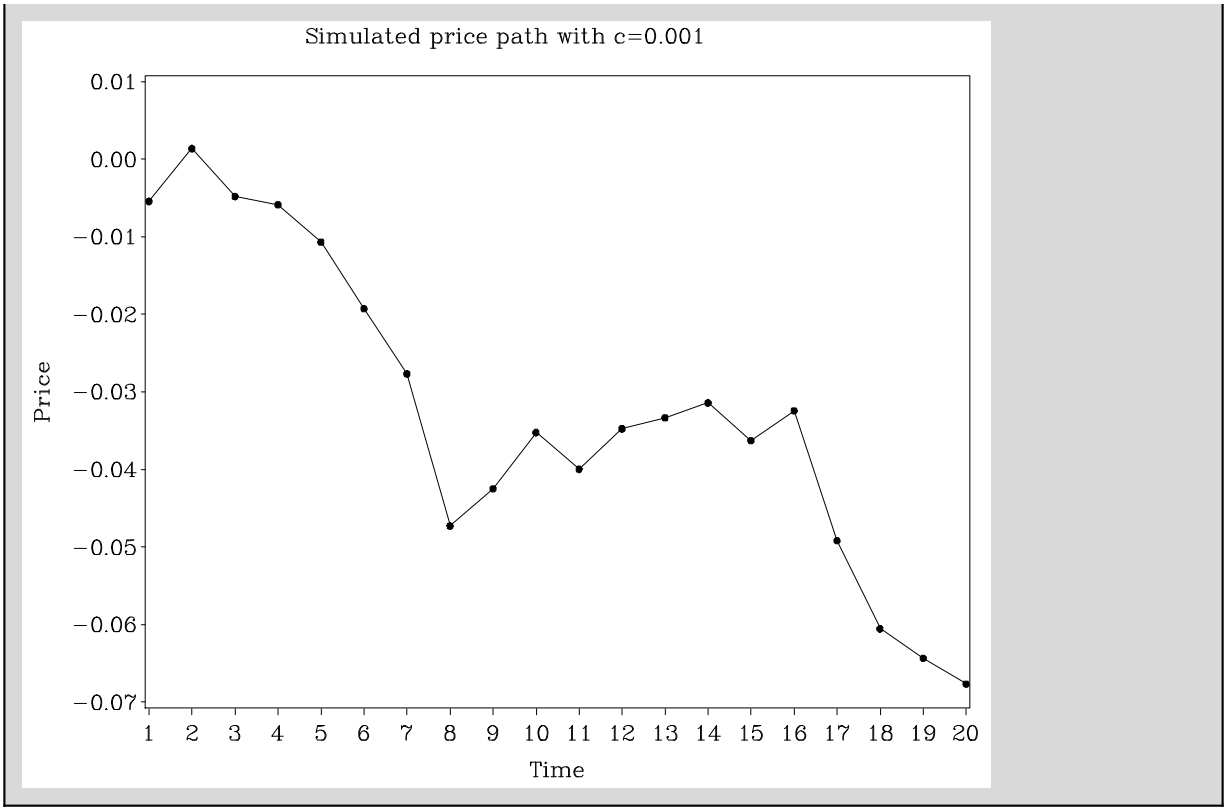

And here are the actual and estimated trade directions.

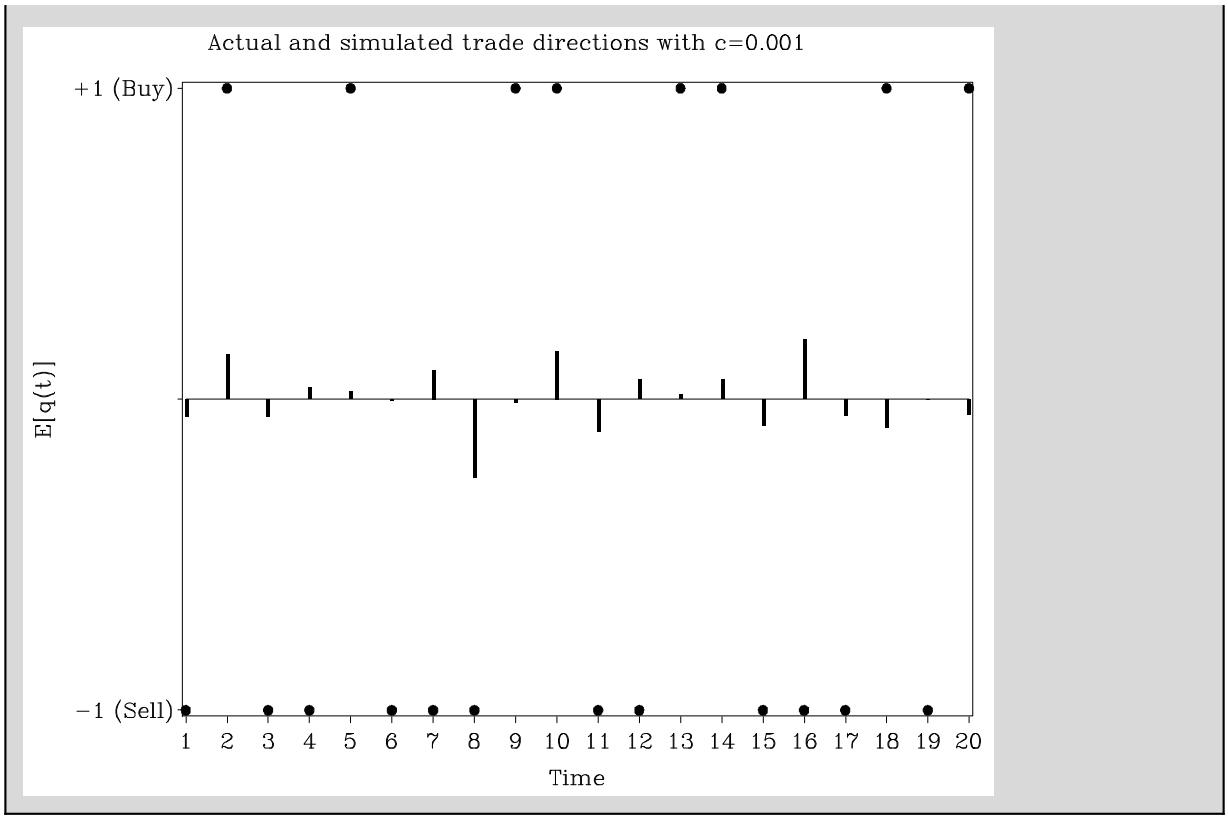

The intuition is as follows. Intuitively, the Gibbs sampler tries to figure out how much of an oberserved price change is transient (due to bid ask effects) and how much is permanent (the efficient price innovation). When `c` is large relative to σ_u, bid-ask bounce generates reversals that are easy to pick out visually, and using the sampler. When `c` is small, though, bid-ask effects are swamped by the innovations in the efficient price, and are not easily discerned.

We'll see that this extends to the parameter estimates as well.

Modification when some of the q's are known.

In some samples, it might happen that the trade directions are known for some subset of the q_t. For these q_t, we don't simulate; we simply leave them at their known values.

This might seem to violate the assumed probability structure of the model in a fundamental way. After all, if the data generating process and our priors are that $q_t = \pm 1$, with equal probability, how can a definite realization be accomodated? The answer is that we're conditioning on the observed data, and the only thing that matters is the prior distribution of the q_t that we don't observe.

By way of a more formal justification, we could assume that the data generating process involves two steps:

First q_t is drawn, ± 1, each with probability $\frac{1}{2}$.

Next, an indicator variable O_t is drawn. With probability π, $O_t = 1$, and the actual q_t is observed. With probability $1 - \pi$, $O_t = 0$, and the actual q_t is unobserved.

As part of the sample, we "observe" the realizations of O_t. That is, we know which q_t are known for sure. If we don't care about modeling the O_t process, letting the observed q_t remain at their known values and simulating the rest corresponds to estimation conditional on the realized O_t. This is a sensible way to proceed.

In doing this, we are implicitly assuming that the o_t process is independent of q_t and u_t. If buys are more likely to be observed than sells, or if the realization of o_t depends on the magnitude of u_t ("Trades are more likely to be observed for large efficient price movements"), then the o_t process is informative, and we may wish to model it explicitly.

Modification when some of the q's are zero.

The U.S. daily CRSP dataset reports closing prices. But if there is no trade on day t, then the reported price is the midpoint of the closing bid and ask. This event is indicated in the data (using a negative price). Essentially, $p_t = m_t$.

Formally, we can incorporate this into the data generating process by noting that $p_t = m_t + c\, q_t = m_t$ when $q_t = 0$.

We can handle this situation in a fashion similar to the known-q_t case discussed above. If the price reported at time t is a quote midpoint, we set $q_t = 0$, and don't resample it during the sweeps.

Formally, this can be justified by letting o_t denote an indicator variable of whether or not there was a trade. Estimation can then proceed condtional on the o_t realizations. Here as well, we're implicitly assuming idependence. We're ruling out (or at least not modeling), for example, the possibility that trades are more likely to occur on days when there are large efficient price changes.

8. Full estimation of the Roll model

Sample runs from *Roll Gibbs 01.sas*

In all cases, the prior on c is $N\left(\mu_c^{\text{Prior}}, \Omega_c^{\text{Prior}}\right)$, restricted to $c > 0$, with $\mu_c^{\text{Prior}} = 0$ and $\Omega_c^{\text{Prior}} = 1$. The prior on σ_u^2 is inverted gamma with $a = b = 1 \times 10^{-6}$.

σ_u = .01; c = 0.01, 100 observations, 20,000 sweeps (first 20% dropped)

Posteriors:

Variable	N	Mean	Dev	Min	Max
SWEEP	16000	12001	4619	4001	20000
SDU	16000	0.0129	0.0016	0.0088	0.0211
C	16000	0.0091	0.0014	0.0001	0.0132

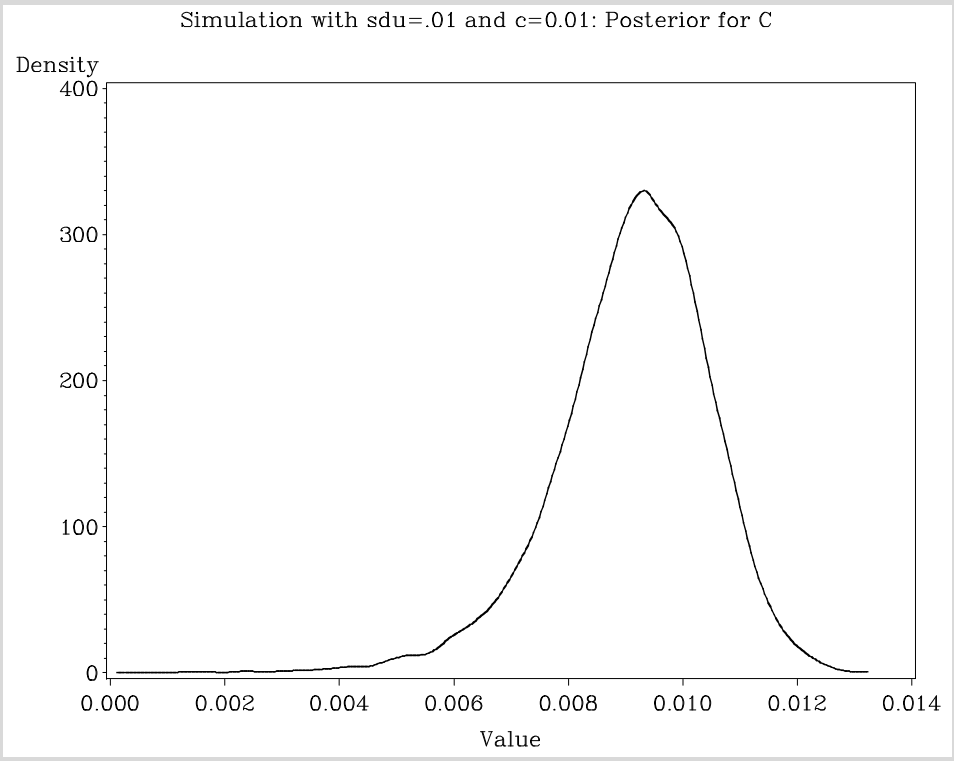

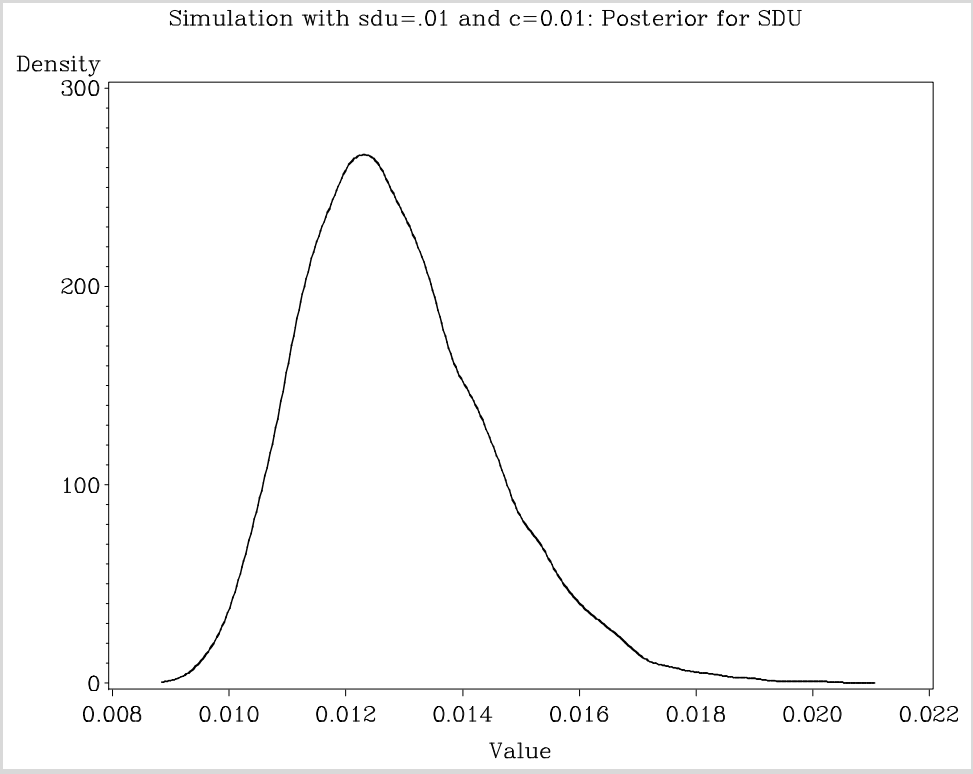

Contour plot of joint posterior:

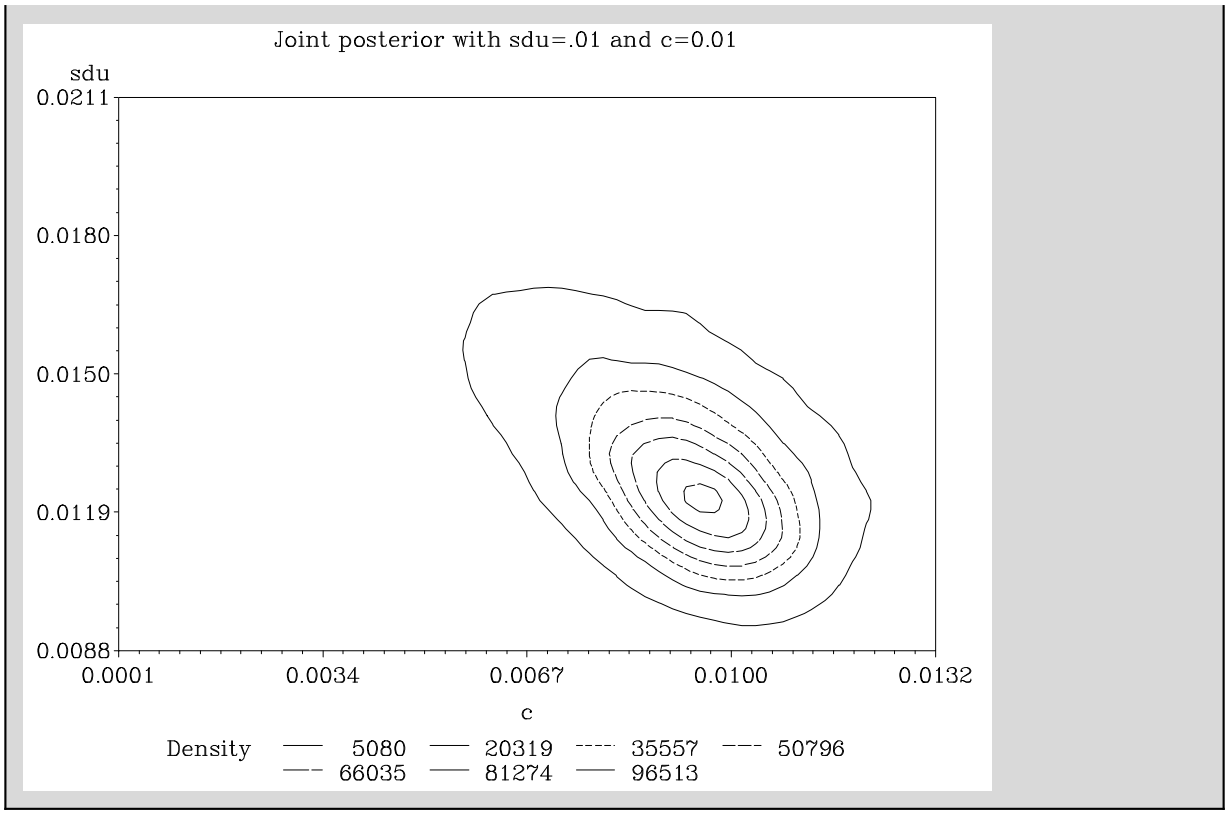

Note the downward slope of posterior. The procedure is trying to allocate volatility, either to the permanent (random-walk) component or to the transient component (effective cost). If more volatility is attributed to the `sdu` component, less is attributed to `c`.

σ_u = .01; c = 0.001, 100 observations, 20,000 sweeps (first 20% dropped)

When `c` << σ_u, the transient cost effects (reversals) are difficult to disentangle from the random-walk component. We still get a reasonably sharp posterior for σ_u (= `sdu`), but the posterior for `c` is broad.

Posteriors:

Variable	N	Mean	Dev	Min	Max
SWEEP	16000	12001	4619	4001	20000
SDU	16000	0.0107	0.0008	0.0082	0.0148
C	16000	0.0014	0.0010	8E-8	0.0057

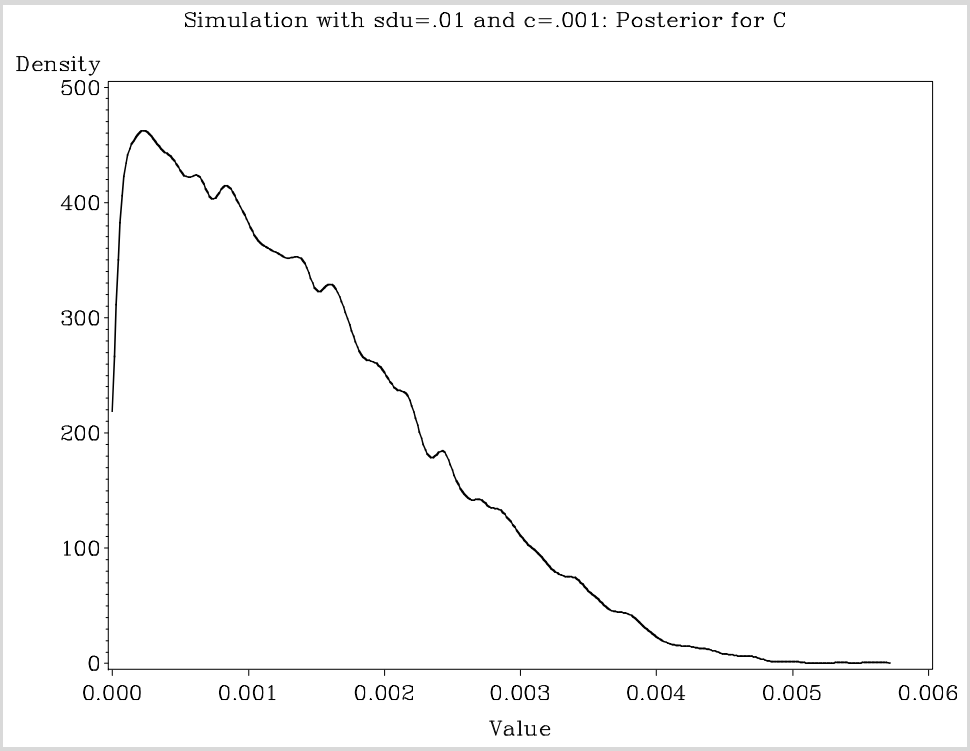

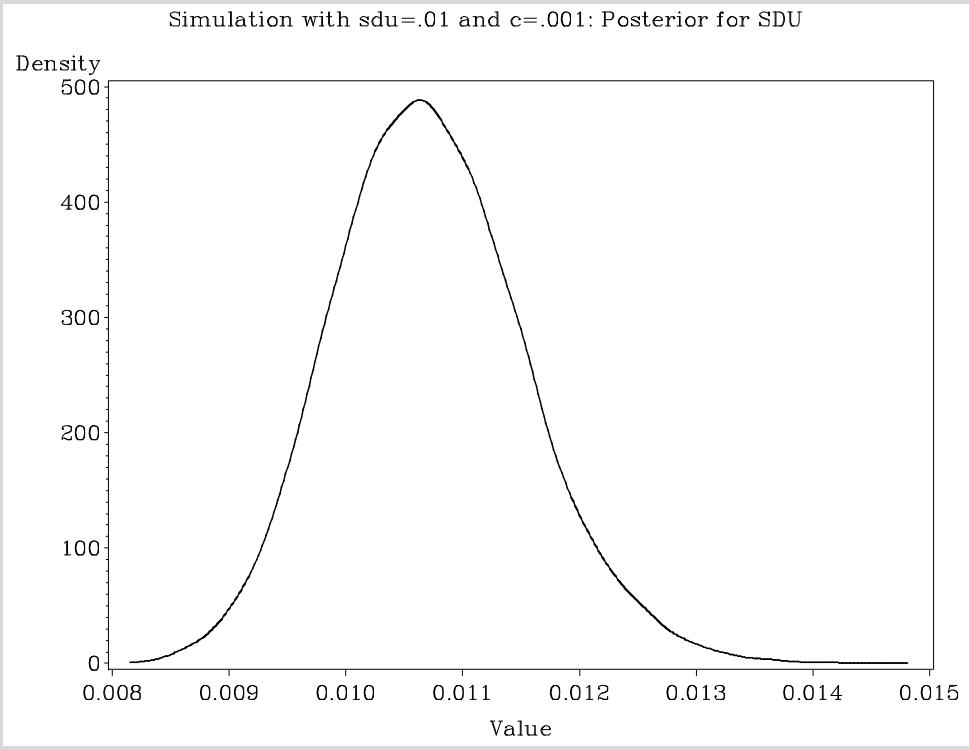

Contour plot of joint posterior:

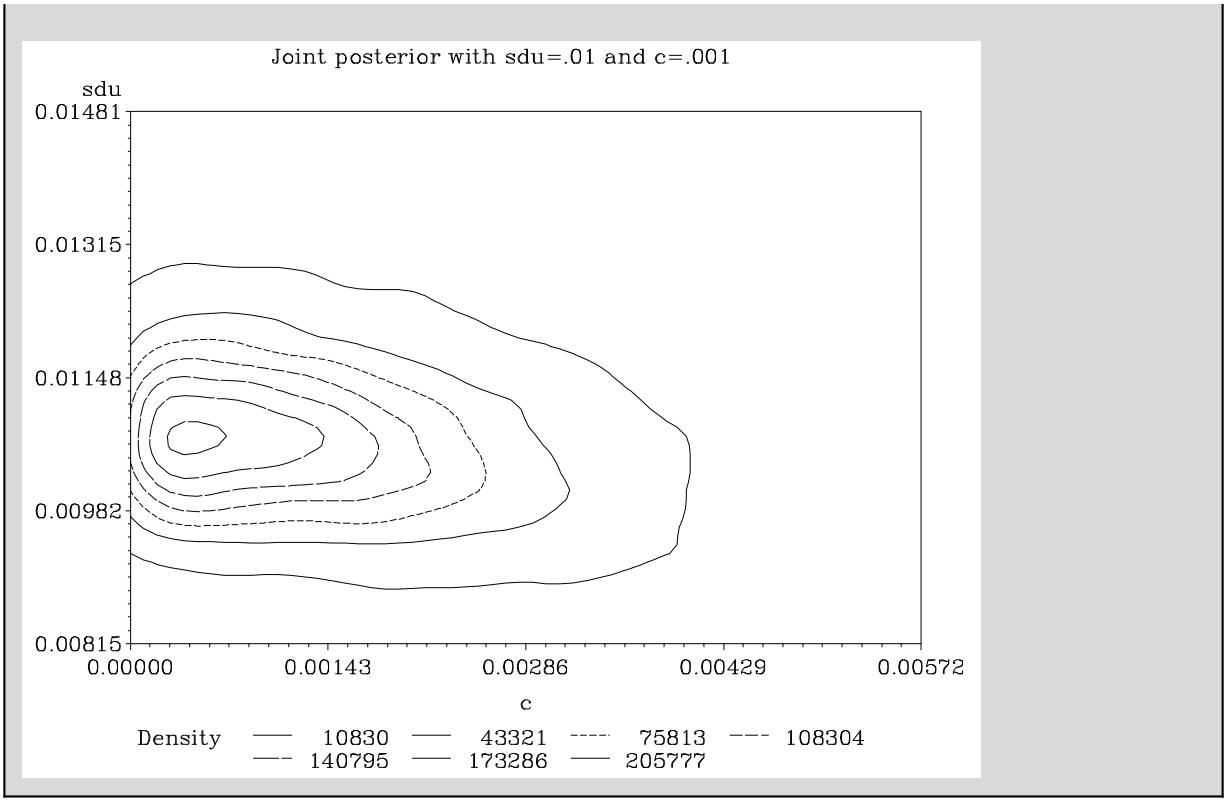

Use of an unrestricted prior for c

Is it really necessary to require $c > 0$? After all, if c really is non-negative, shouldn't the procedure pick it up?

When we run the last problem with no restrictions on c, two things happen.

For all of the q_t, the sample draws average out to zero. (We can't tell whether a given trade is a buy or a sell.)

The posterior for c is bimodal, and symmetric about zero.

Here's an example. 30 observations were simulated with $\sigma_u = 0.01$ and $c = 0.01$. The prior for c was not restricted to be non-negative. Here is the posterior:

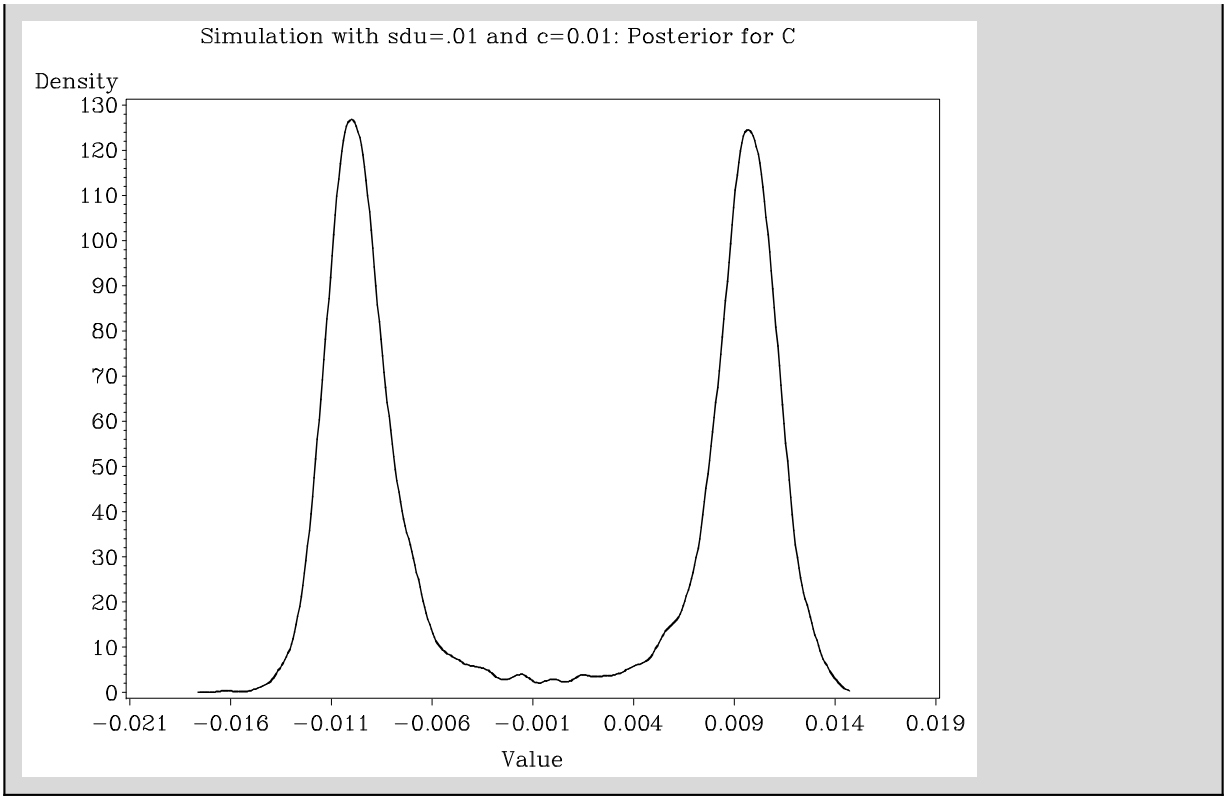

What's happening is this. The Roll model with `c > 0` is observationally equivalent to one in which `c < 0`, and the trade signs are reversed. Without the nonnegativity restriction on `c`, the posterior (and the sampler) will span both possibilities.

Using vague prior for c

In the preceding simulations, the prior for `c` is $N(0, 1)$, restricted to $c > 0$. This is fairly flat over the usual region of interest. (For a US equity, even extreme values of `c`, estimated from trade and quote data, are rarely above 0.05.) Nevertheless, there is some curvature. Why not remove all doubt and set the prior to, say, $N(0, 1\,000\,000)$?

The problem with this is that under some circumstances we may need to make a draw from the prior. This is not common, but in a small sample, with $c << \sigma_u$, over many draws, the following situation may arise. Suppose that on a particular sweep, the trade direction indicators are drawn to have the same sign: $q_1 = \ldots = q_T = +1$ or $q_1 = \ldots = q_T = -1$. In this case, all of the Δq_t are zero. But the Δq_t are the r.h.s. variables in the price change regression. If an explanatory variable in a regression has no variation, the regression is completely uninformative. In this case, the draw of c for that sweep must be made from the prior.

9. Return factors

Return factors are logically introduced by adding them to the efficient price change specification, e.g.,

$$m_t = m_{t-1} + f_t' b + u_t$$

where f_t is a $K \times 1$ vector of known factor realizations and `b` is the vector of loadings (factor coefficients). The factor terms then appear in the trade price specification:

$$\Delta p_t = c\,\Delta q_t + f_t' b + u_t$$

For example, a market model for stock `i` might be specified as:

$$\Delta p_{it} = c_i \Delta q_{it} + b_i r_m + u_t$$

where r_m is the return on a market index.

The *b* are estimated as another coefficient in the Bayesian regression (the same regression in which c is estimated). The parameters of the coefficient prior μ_β^{Prior} and $\Omega_\beta^{\text{Prior}}$ are expanded to include the new coefficients, and they are drawn (simulated) in the same step as the draw for c.

The addition of a return factor will generally increase the resolution between permanent (random-walk) and transient (bid-ask) components. In practice, this accomplishes two things:

We'll get a better estimate of c.

We'll also generally get better estimates of b.

The first of these is pretty straightforward: the explanatory power of the factors reduces the residual variance. The second point may require some expansion. bs are conventionally estimated using daily price changes. These daily price changes are contaminated by bid-ask bounce. For some stocks, bid-ask bounce may be large compared to the factor-induced and idiosyncratic changes in the efficient price, leading to large estimation errors in the bs. Estimating a specification that includes a $c \triangle q_t$ term effectively allows us to estimate bs on price series that are purged of the bid-ask bounce.

Market model example (*RollGibbsBeta 01.sas)*

The specification is the one-factor model described above The parameters are:
$c = 0.01;\ r_m \sim N\left(0,\ \sigma_m^2\right)$ with $\sigma_m = 0.01;\ \sigma_u = 0.01;\ \beta = 1.1.$
100 observations were simulated; 10,000 sweeps (first 20% discarded).
The posterior summary statistics are:

Variable	N	Mean	Std Dev	Minimum	Maximum
SWEEP	8000	6001	2310	2001	10000
SDU	8000	0.0106	0.0012	0.0075	0.0173
C	8000	0.0109	0.0009	0.0062	0.0140
BETA	8000	1.098	0.140	0.578	1.740

For purposes of comparison, β was also estimated in the usual way (OLS). The estimated specification:

$$r = \underset{(0.186)}{1.132}\, r_m + e.$$

Here are the posteriors:

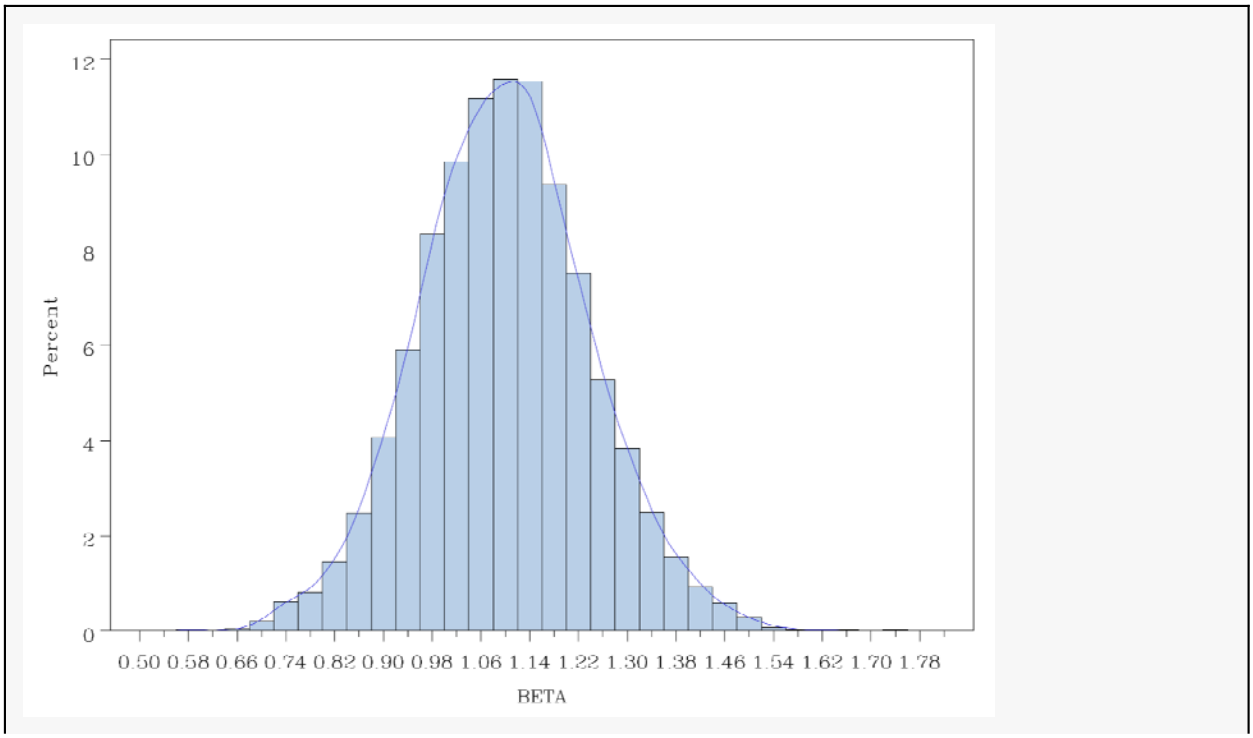
Percent
12
10
8
6
4
2
0
0.50 0.58 0.66 0.74 0.82 0.90 0.98 1.06 1.14 1.22 1.30 1.38 1.46 1.54 1.62 1.70 1.78
BETA

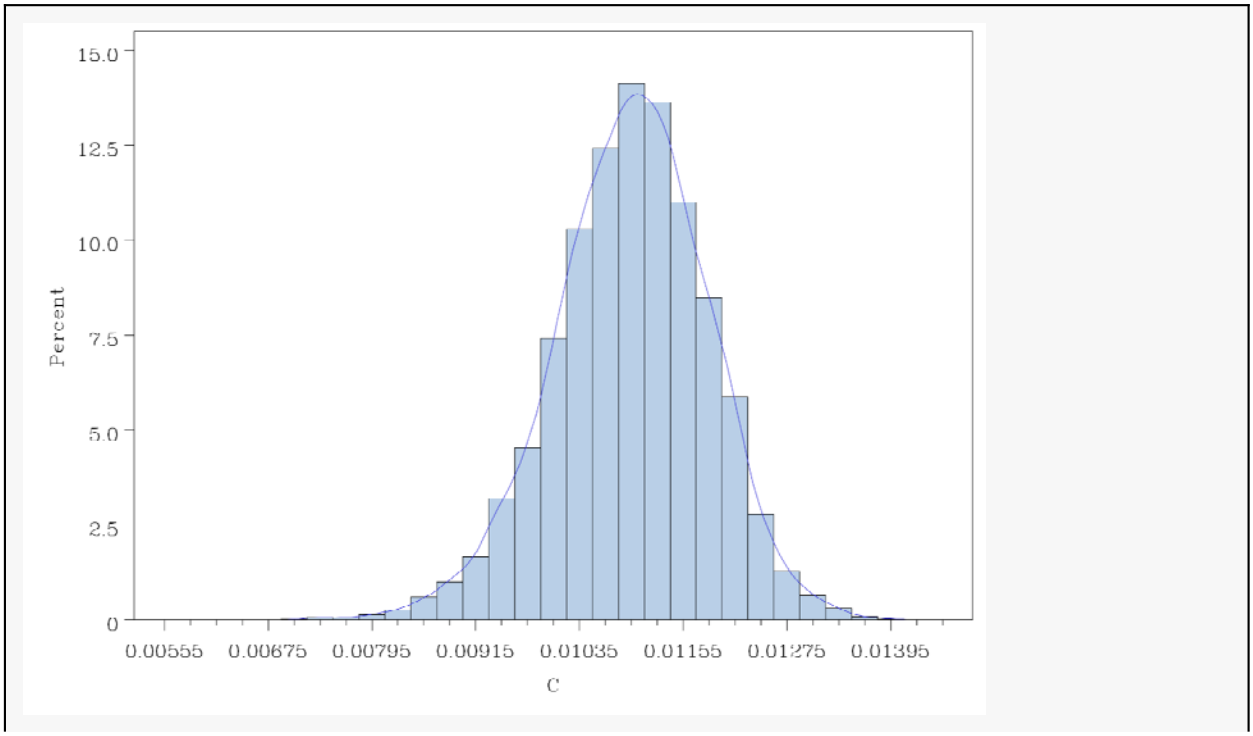
Percent
15.0
12.5
10.0
7.5
5.0
2.5
0
0.00555 0.00675 0.00795 0.00915 0.01035 0.01155 0.01275 0.01395
C

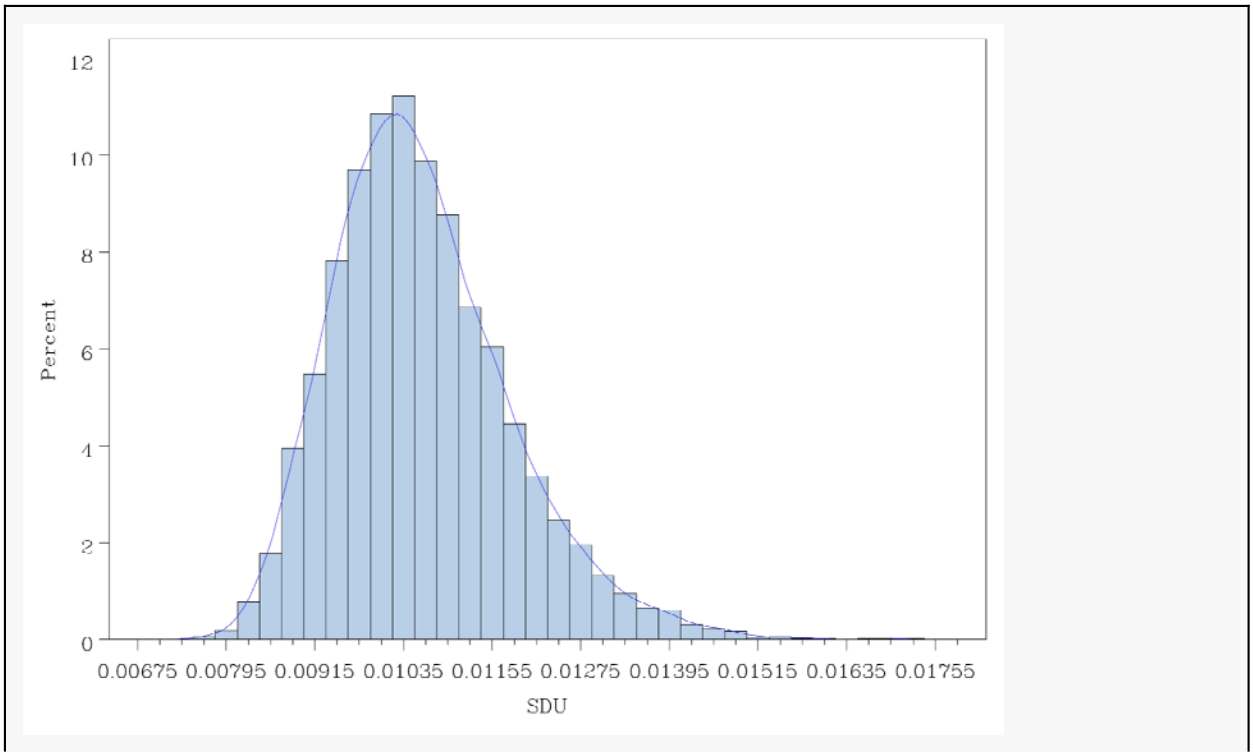

Extensions

Variation in *c*

Usual way of analyzing variation in a liquidity proxy:

1. Estimate liquidity proxy over short subsamples.
2. Use the estimates as dependent variables in a subsequent regression.

E.g., Pastor and Stambaugh construct monthly γ estimates where each estimate is based on all the trading days in the month, and then estimate liquidity betas from a portfolio variant of $\hat{\gamma}_{it} = \alpha_i + \beta_i \hat{\gamma}_{mt} + e_{it}$

Small-sample Gibbs estimates of *c* don't have nice properties. They are biased. In a small sample, the posterior is dominated by the prior. The prior for *c* is typically half-normal:

```
Plot[PDF[NormalDistribution[0, .10], x], {x, 0, .3}, AxesOrigin → {0, 0}, PlotRange → All]
```

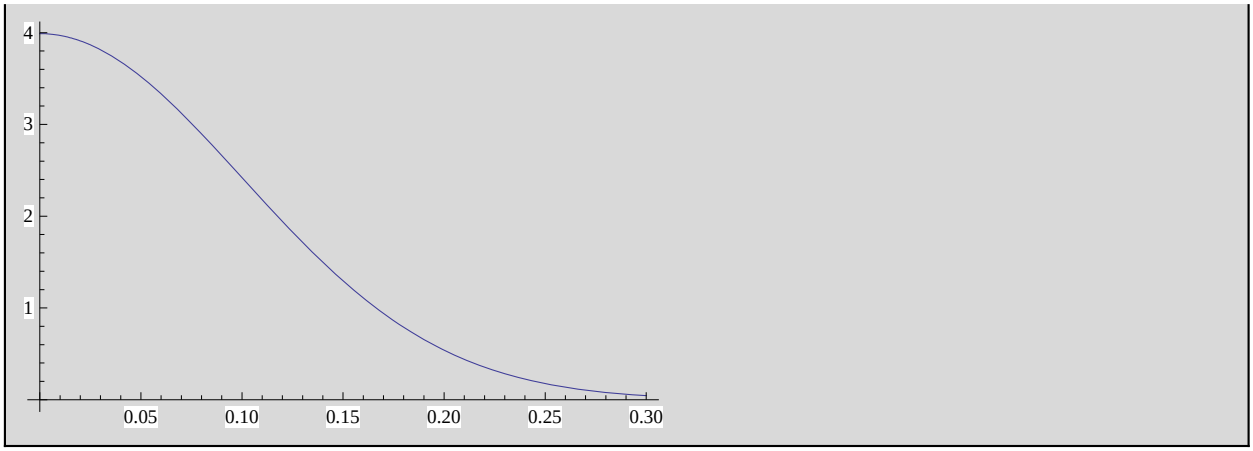


```
ExpectedValue[Abs[x], NormalDistribution[0, .1], x]
```

```
0.0797885
```

It is better to model variation in c directly, incorporating the functional form of c into the Gibbs estimation. Here are some scenarios.

Time-series variation

$m_t = m_{t-1} + u_t$
$p_t = m_t + c_t\, q_t$ where $c_t = z_t\, \gamma$, i.e., a linear function of observed variables z_t

Parameter and/or data restrictions may be necessary to ensure $c_t > 0$.

The Gibbs sampler can be modified to include the simulation of γ as an additional step.

$$\Delta p_t = c_t\, q_t - c_{t-1}\, q_{t-1} + u_t \;=\; (q_t\, z_t - q_{t-1}\, z_{t-1})\, \gamma + u_t$$

Given ("conditioning on") the qs and zs, this is a normal regresssion model with coefficient vector γ.
z_t might include indicators for event days, etc. ("Are trading costs higher on earnings announcement days?")

Cross-sectional variation

$m_{it} = m_{i,t-1} + u_{it}$
$p_{it} = m_{it} + c_{it}\, q_{it}$ where $c_{it} = z_{it}\, \gamma_i$

Here, z may include cross-sectional variables (e.g., market capitalization) etc.
If the u_{it} (or q_{it}) are cross-sectionally dependent, the $\Delta\, p_t$ regressions should be estimated as a SUR system.
We can avoid this if we impound all cross-sectional dependence in a market factor:

$$m_{it} = m_{i,t-1} + \beta_i\, M_{it} + u_{it}$$

Latent variation

$p_{it} = m_{it} + c_{it}\, q_{it}$ where $c_{it} = \tilde{z}_t\, \gamma_i$

Here, $\tilde{z}_t$ is interpreted as an unobserved (latent) liquidity factor that is common to all stocks at given time. We can estimate it because we can write the price change as:

$$\Delta p_{it} = c_{it}\, q_{it} - c_{it-1}\, q_{it-1} + u_{it} = \tilde{z}_t(\gamma_i\, q_{it}) - \tilde{z}_{t-1}(\gamma_i\, q_{it-1}) + u_{it}$$

Written in this fashion, conditional on everything else, $\tilde{z}_t$ is the coefficient in a cross-sectional regression.

... and more

Bayesian estimation of asset pricing models is well-developed. MCMC estimators (like the Gibbs) are easy to implement because they are modular. Thus, we could estimate an asset pricing model directly on the *m*s, i.e., returns purged of bid-ask bounce.

Stochastic volatility models are also easily estimated in the Bayesian framework. An SV model could be grafted on the Roll setup, possibly allowing for a link between volatility and the trading cost.

Stochastic duration models may be useful because information intensity is often modeled as being inversely related to intra-trade or intra-event duration.

Missing data models have been used in asset pricing studies to address the problem that we rarely have a full panel of asset returns over time.

A neat Bayesian trick (aside)

Some of the models described above call for panel estimations. Sometimes there are short-cuts.

Suppose that we have a panel regression $y_{it} = X_{it}\,\beta + e_{it}$ where the e_{it} are not cross-sectionally correlated $\left(E e_i\, e'_j = 0 \text{ for } i \neq j\right)$. We could estimate this as a panel system:

$$\begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{pmatrix} = \begin{pmatrix} X_1 \\ X_2 \\ \vdots \\ X_N \end{pmatrix} \beta + \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_3 \end{pmatrix}$$

where y_1 is the column vector $(\; y_{1,1} \quad y_{1,2} \quad \cdots \quad y_{1,T} \;)'$, etc.

Alternatively, suppose that the prior for β is $N(\mu_0,\, \Omega_0)$. We proceed sequentially as follows.

- Estimate $y_1 = X_1\,\beta + e_1$, obtaining a posterior for β, denoted $N(\mu_1,\, \Omega_1)$.
- Use $N(\mu_1,\, \Omega_1)$ as a prior in estimating $y_2 = X_2\,\beta + e_2$, obtaining a posterior for β, denoted $N(\mu_2,\, \Omega_2)$
- Use $N(\mu_2,\, \Omega_2)$ as a prior in estimating $y_3 = X_{23}\,\beta + e_3$, obtaining a posterior for β, denoted $N(\mu_3,\, \Omega_3)$

... and so on. It is sometimes easier to loop in this fashion, rather than lay out and set up the large matrices necessary to run the SUR.

Further note: you can also form a posterior for a single regression by updating observation-by-observation. This is essentially what a Kalman filter does (sequentially, in a time-series context).

Asymmetric information

The Gibbs sampler has been applied to the classic Roll model, in which the increments in the efficient price, u_t, are assumed to be independent of the trade direction indicators, q_t.

Could we roughly check this independence by estimating the correlation between the u_t and q_t (draw by draw)? Remember that the estimates are coming from the regression

$$\Delta p_t = c\,\Delta q_t + u_t$$

In a regression, the estimated residuals are uncorrelated with the explanatory variables by construction.

A generalized Roll model that does allow for the trade effects associated with asymmetric information is:

$$m_t = m_{t-1} + \lambda q_t + u_t$$
$$p_t = m_t + c q_t$$

This implies

$$\Delta p_t = (\lambda + c)\, q_t + c q_{t-1} + u_t$$

The q draws for this model are slightly more involved. See Hasbrouck (2004).

There is a limitation here. The original Roll model holds its form under time aggregation. For example:

$$p_t - p_{t-2} = c(q_t - q_{t-2}) + u_t + u_{t-1}$$

That is, the bid-ask bounce term is determined by the qs are the endpoints of the observation interval, and the aggregate disturbance has the same form as the original disturbances.

In the generalized Roll model, though:

$$p_t - p_{t-2} = c(q_t - q_{t-2}) + \lambda(q_t + q_{t-1}) + u_t + u_{t-1}$$

This involves not just the qs at the endpoints, but also the intermediate value q_{t-1}. If we were applying the model to closing trade prices, the daily price change between day $d-1$ and day d would look like:

$$p_d - p_{d-1} = c(q_N - q_1) + \lambda(q_1 + q_2 + \cdots + q_N) + (u_1 + u_2 + \cdots + u_N)$$

where the subscripts on the q and u run over the N trades on day d. That is, the sum $(q_1 + q_2 + \cdots + q_N)$ is the net buy order flow over the day. In principle we could compute the conditional distribution for this sum, and make a Gibbs draw. In practice, though, most of the time we don't know N. This complicates matters enormously.

Volume

Suppose that we're in a Kyle world and y_i denotes the net signed volume in the ith market. That is:

$$\Delta p_i = \lambda y_i \text{ where } y_i \sim N(0, \sigma_y^2)$$

If we had a single observation on Δp and y, we could estimate λ. It really wouldn't be necessary to know *signed* volume, though. It would suffice to know the *unsigned* volume $|y_i|$, since $|\Delta p_i| = \lambda |y_i|$.

This intuition, that we can estimate λ from the absolute values, underlies the Amihud illiquidity measure.

Note: in the limit of the Kyle model, where aggregate order flow y is a Wiener process, we might be tempted to interpret $|dy|$ as "volume". This runs into the problem that a Wiener process has unbounded variation.

Aggregation

Things are more complicated under aggregation. Usually we know a day's total trading volume

$$V = |y_1| + |y_2| + \cdots + |y_N|$$

What is $f(y_1 + y_2 + \ldots + y_N \mid V)$? (And we don't know N.)

Note: If $y_i \sim N(0, \sigma_y^2)$, then $|y_i| / \sigma_y$ is distributed as a χ variate with one degree of freedom. Note: *not* a χ^2 variate. This is unfortunate because χ^2 variates aggregate nicely, while χ variates do not.

Possible approaches.

- Assume that all trades are for 100 shares. Then we're back to inferring $q_1 + q_2 + \ldots + q_N$. This may be easier.

Easley et al. assume that buys arrive with a Poisson intensity λ_B and sells with a Poisson intensity λ_S. From the aggregation properties of Poisson processes, the trade arrival process has intensity $\lambda_B + \lambda_S$. This implies a distribution of volume. Furthermore, the net number of buys is unconditionally distributed as a Skellam variate. What we need, though is the condition distribution of the net number of buys (conditional on the number of trades).

When the time-aggregation period is sufficiently long that we have a large number of buys and sells, it may be useful to apply the normal approximation to the Poisson distributions. Even when the time period is short, we might still use the normal approximation to generate candidate draws for a Metropolis-Hastings algorithm.

Public (non-trade) information

$$\Delta p_i = \lambda y_i + u_i \Rightarrow \frac{|\Delta p_i|}{|y_i|} = \frac{|\lambda y_i + u_i|}{|y_i|} \leq \frac{\lambda|y_i| + |u_i|}{|y_i|} = \lambda + \frac{|u_i|}{|y_i|}$$

This suggests that I will be a downward-biased estimate of λ.

And finally: a caution

The inferences that we draw from a statistical model depend on the model and the data. We can get stronger inferences from using a tighter model or getting more data (not just more observations, but more kinds of variables).

Sometimes we require data that is either inherently latent or is simply unobtainable for practical reasons. MCMC methods allow us to estimate these data, but only if we are willing to specify a stronger stronger model. That's the trade-off.

The caution is that it is very easy to specify a fanciful model loaded with latent data. If the mixing and convergence properties of the simulation estimator are poor, the model can appear to be well-identified.

Addendum: Transformations of the Gamma distribution

When `a = n / 2` and `b = 1 / 2`, the Gamma becomes:

```
PDF[ChiSquareDistribution[n], x]
```

$$\frac{2^{-n/2}\, e^{-x/2}\, x^{-1+\frac{n}{2}}}{\text{Gamma}\left[\frac{n}{2}\right]}$$

We sometimes need to determine the pdf's of σ_u^2 and/or σ_u.

If $\frac{1}{\sigma_u^2}$ is Gamma, then what is the pdf of σ_u^2?

Recall that if $y = g(x)$, then $f(y) = f\left(g^{-1}(y)\right)\left|g^{-1\prime}(y)\right|$

```
gRule = {g[x_] :> x^-1, gi[y_] :> y^-1};
```

```
fy =
 PDF[GammaDistribution[a, λ], x] * D[-gi[y] /. gRule, y] /. λ → 1 / b /. x → gi[y] /. gRule //
  Simplify
```

$$\frac{\left(\frac{1}{b}\right)^{-a} e^{-\frac{b}{y}} \left(\frac{1}{y}\right)^{1+a}}{\text{Gamma}[a]}$$

Verify that this integrates to unity:

```
Integrate[fy, {y, 0, ∞}, Assumptions -> {a ∈ Reals, b ∈ Reals, a > 0, b > 0}]
```

```
1
```

Compute the expectation:

```
Integrate[y fy, {y, 0, ∞}, Assumptions -> {a ∈ Reals, b ∈ Reals, a > 1, b > 0}]
```

$$\frac{b}{-1+a}$$

If $\frac{1}{\sigma_u^2}$ is Gamma, then what is the pdf of σ_u (the standard deviation)?

```
gRuleSD = {g[x_] :> x^(-1/2), gi[y_] :> y^-2};
```

```
fy = PDF[GammaDistribution[a, λ], x] * D[-gi[y] /. gRuleSD, y] /. λ → 1 / b /. x → gi[y] /.
    gRuleSD // Simplify
```

$$\frac{2\left(\frac{1}{b}\right)^{-a} e^{-\frac{b}{y^2}} \left(\frac{1}{y^2}\right)^{a}}{y\,\text{Gamma}[a]}$$

Verify that this integrates to unity:

```
Integrate[fy, {y, 0, ∞}, Assumptions -> {a ∈ Reals, b ∈ Reals, a > 0, b > 0}]
```

```
1
```

Compute the expectation:

```
Integrate[y fy, {y, 0, ∞}, Assumptions -> {a ∈ Reals, b ∈ Reals, a > 1, b > 0}]
```

$$\frac{\sqrt{b}\ \text{Gamma}\left[-\frac{1}{2}+a\right]}{\text{Gamma}[a]}$$

References

Ball, Clifford A. and Tarun Chordia, 2001, True spreads and equilibrium prices. Journal of Finance 56, 1801-35.

Carlin, Bradley P. and Thomas A. Louis, 2000, Bayes and Empirical Bayes Methods for Data Analysis. London, Chapman and Hall.

Geweke, John, 2005, Contemporary Bayesian Statistics and Econometrics. New York, John Wiley and Sons.

Hasbrouck, Joel, 1999, Security bid/ask dynamics with discreteness and clustering: Simple strategies for modeling and estimation. Journal of Financial Markets 2, 1-28; doi:10.1016/S1386-4181(98)00008-1

Hasbrouck, Joel, 2004, Liquidity in the futures pits: Inferring market dynamics from incomplete data. Journal of Financial and Quantitative Analysis 39.

Hasbrouck, J. (2009). Trading costs and returns for US equities: estimating effective costs from daily data. Journal of Finance 64(3): 1445-1477 DOI link: http://dx.doi.org/10.1111/j.1540-6261.2009.01469.x

Kim, Chang-Jin and Charles R. Nelson, 2000, State-space models with regime switching. Cambridge, Massachusetts, MIT Press.

Lancaster, Tony, 2004, An Introduction to Modern Bayesian Econometrics. Malden (MA), Blackwell Publishing.

Shephard, Neil, 2005, Stochastic Volatility. Oxford, Oxford University Press.

Tanner, Martin A., 1996, Tools for Statistical Inference. New York, Springer-Verlag.