Some Linguistics of Quantitative Finance

Pat Burns
2016 May
R in Finance, Chicago

I started with a logic proposition:

- I'd rather not get through all my slides
- It's almost Saturday night (the talk was Saturday after lunch)

therefore

Heckling is encouraged

(However in the event the audience failed to heckle.)

Agenda

- Increase your ignorance
- Start a quant community to-do list

(Suggested heckle: What if I don’t want to be ignorant? And I don’t.

Response: This is **my** agenda, yours may vary.)

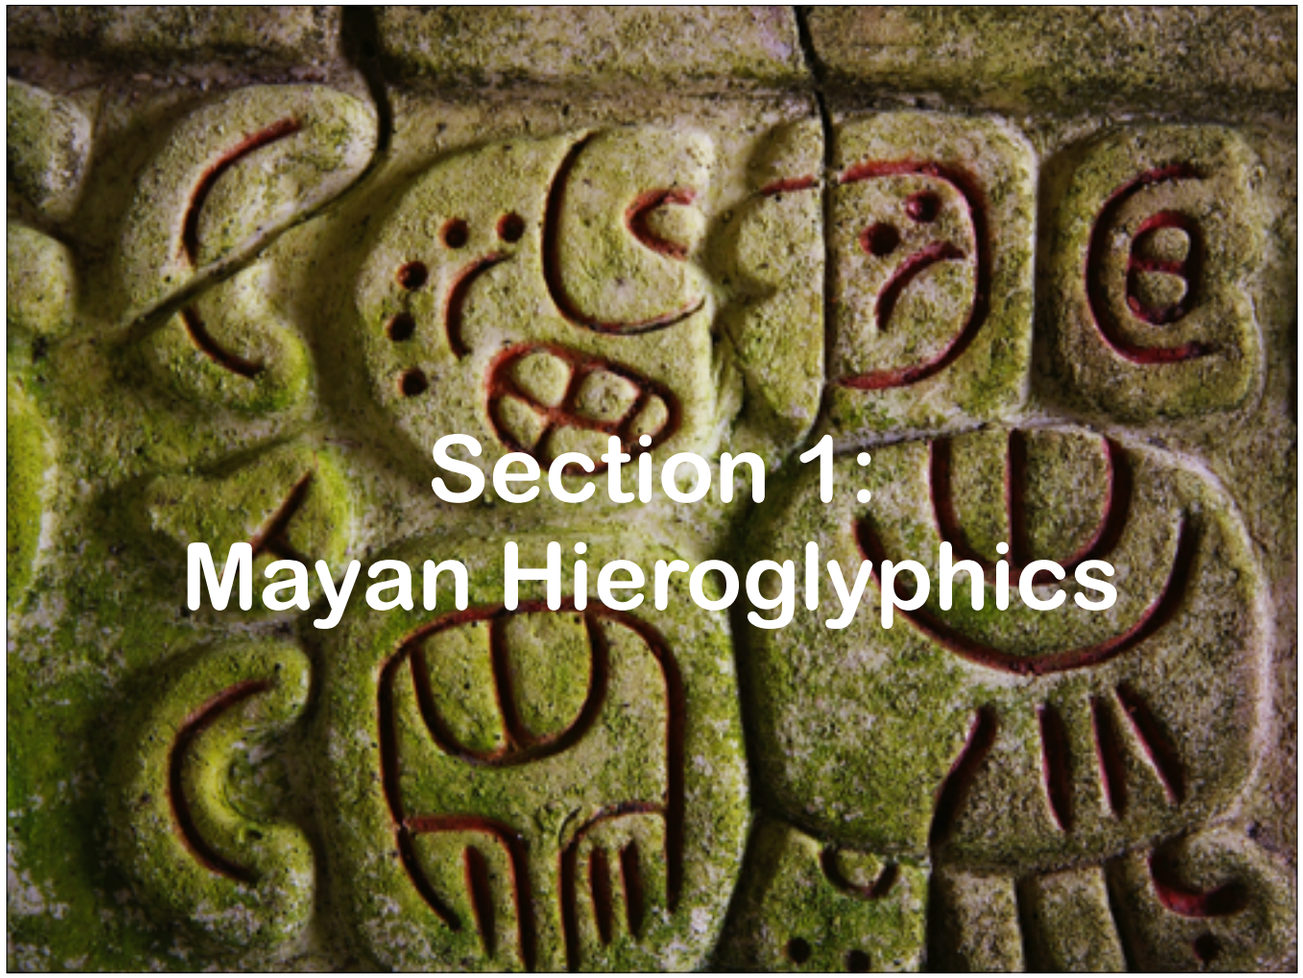

Let’s jump into the linguistics.

Photo from istock.

Benjamin Whorf

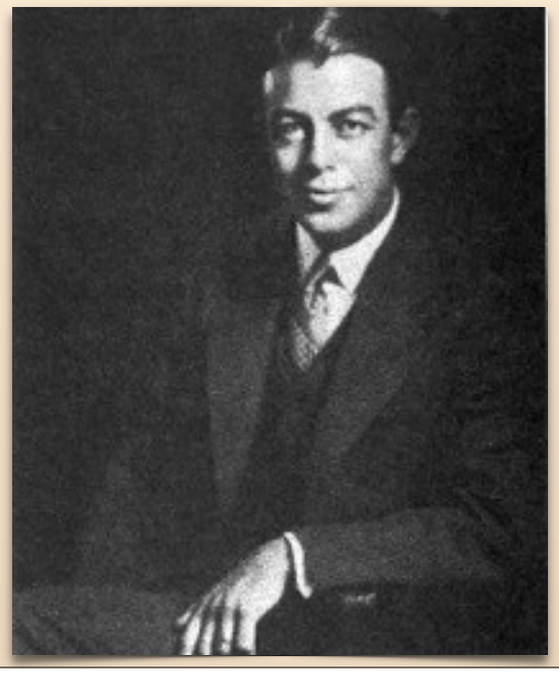

Whorf was a risk manager before there was risk management. He worked for the Hartford (Wallace Stevens worked there at the same time, but there is no record of them meeting).

Photo from wikipedia.

Whorf’s job was to shave the dice for fire insurance.

But that was just his day job.

Photo from istock.

He was really a linguist.

These two streams mingled together, and Whorf noticed that an awful lot of fires were started by …

empty gasoline cans.

But empty gas cans aren't really empty. They are full of gasoline vapor — that's the explosive part. (Hydrocarbons are surrounded by oxygen molecules.)

This is an example of …

Photo from istock.

Sapir-Whorf Hypothesis

Our language controls the way we think

(Paul Teetor's talk approximately 67 minutes later was explaining an instance of the Sapir-Whorf hypothesis.)

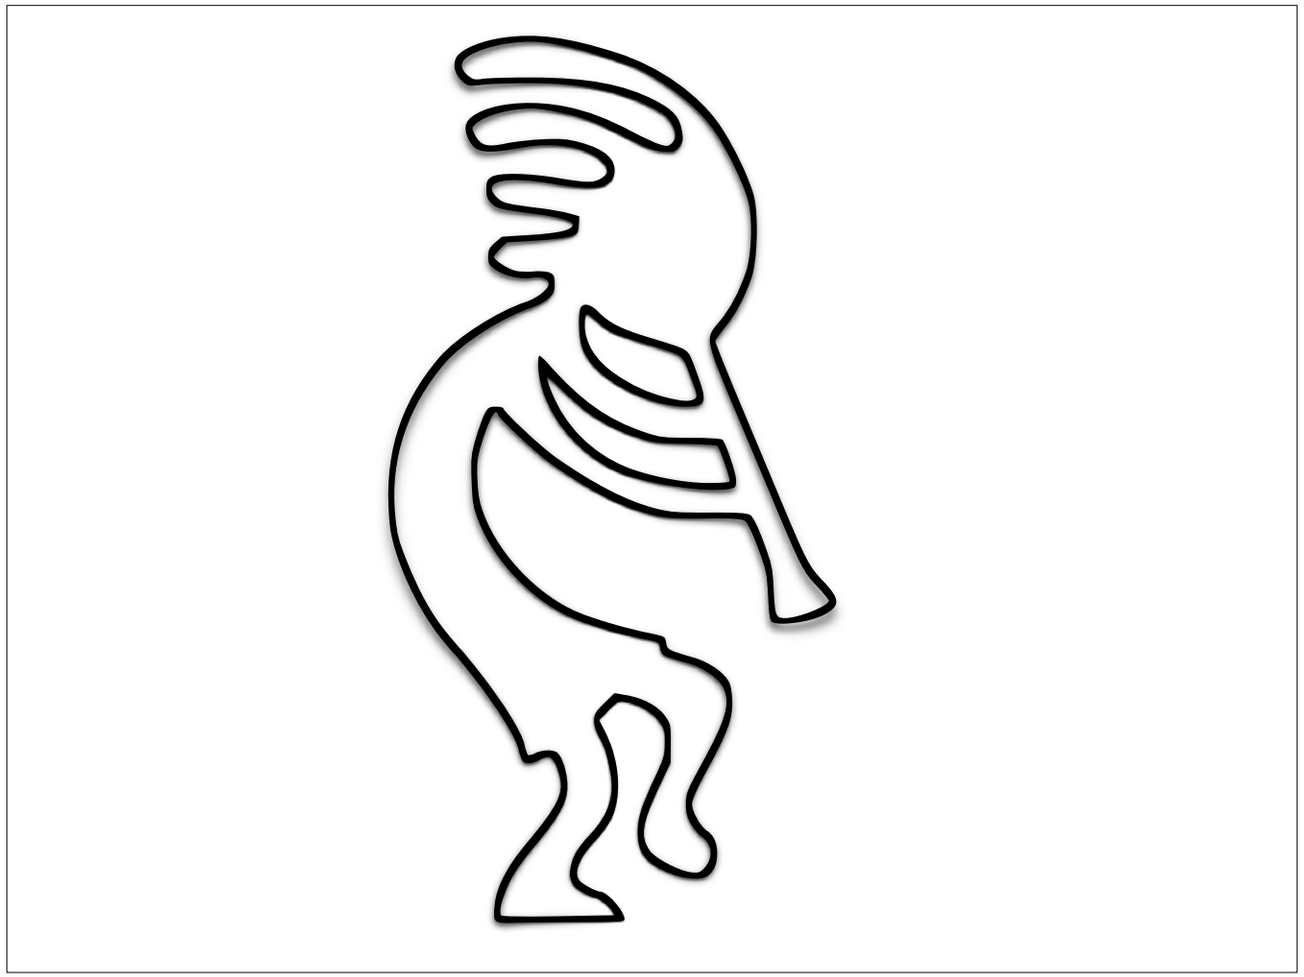

One of the languages that Whorf studied was Hopi.

(This is Kokopelli who plays the flute and is a god of fertility and the harvest, and a prankster.)

Image from istock.

In european languages we say that nouns are things and verbs are actions, with a bit of fine print.

(The more you examine this the more holes you find.)

Photo from istock.

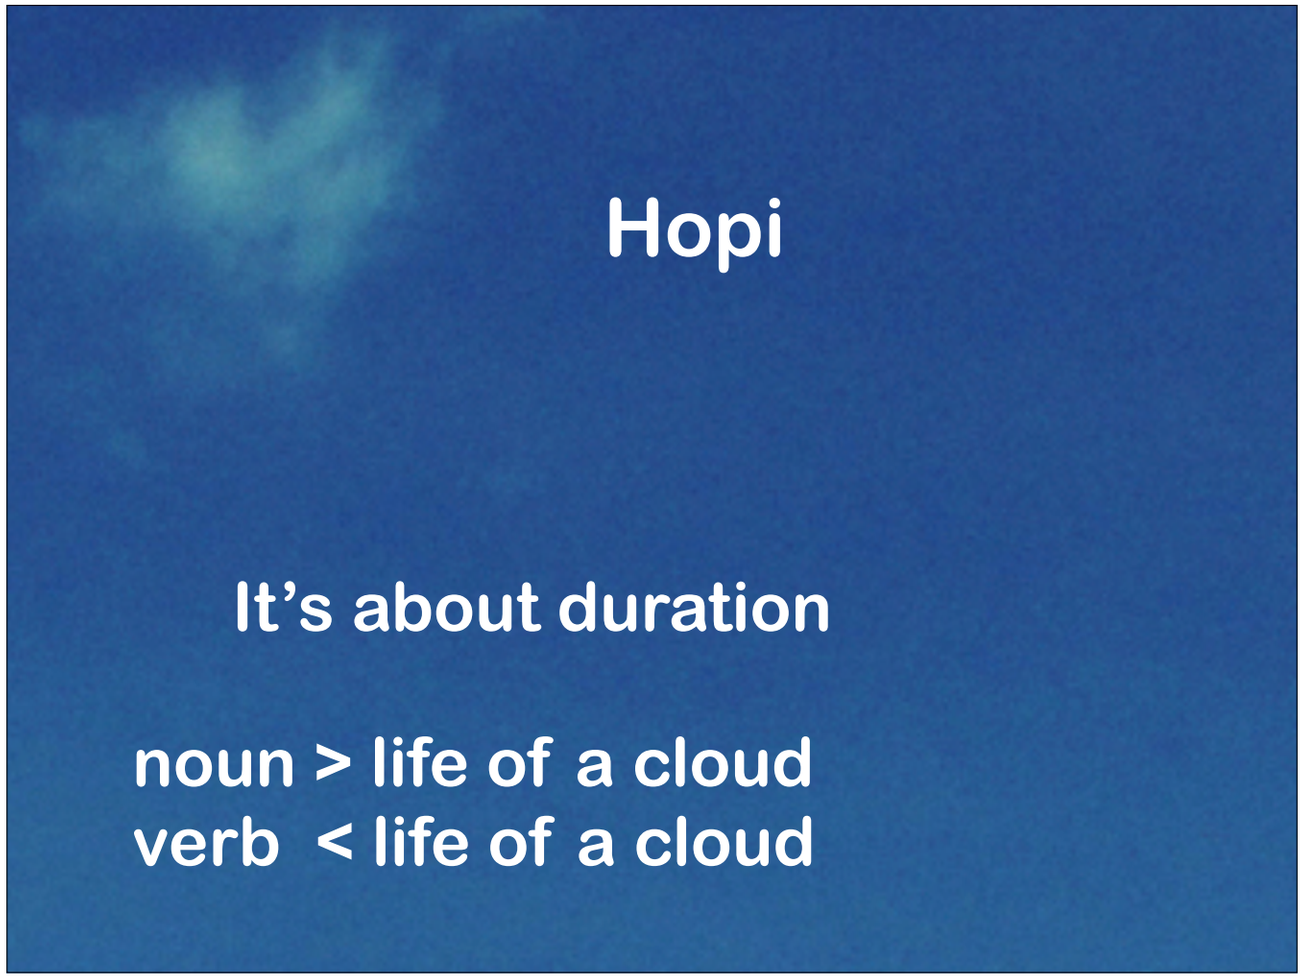

Whorf said that Hopi is different.

Nouns last a long time, verbs last a short time. The boundary between the two is about the length of time that a cloud lasts in the Southwest.

(A verb is conjugated with tenses, there’s no need for tenses if there is permanence.)

Photo from istock.

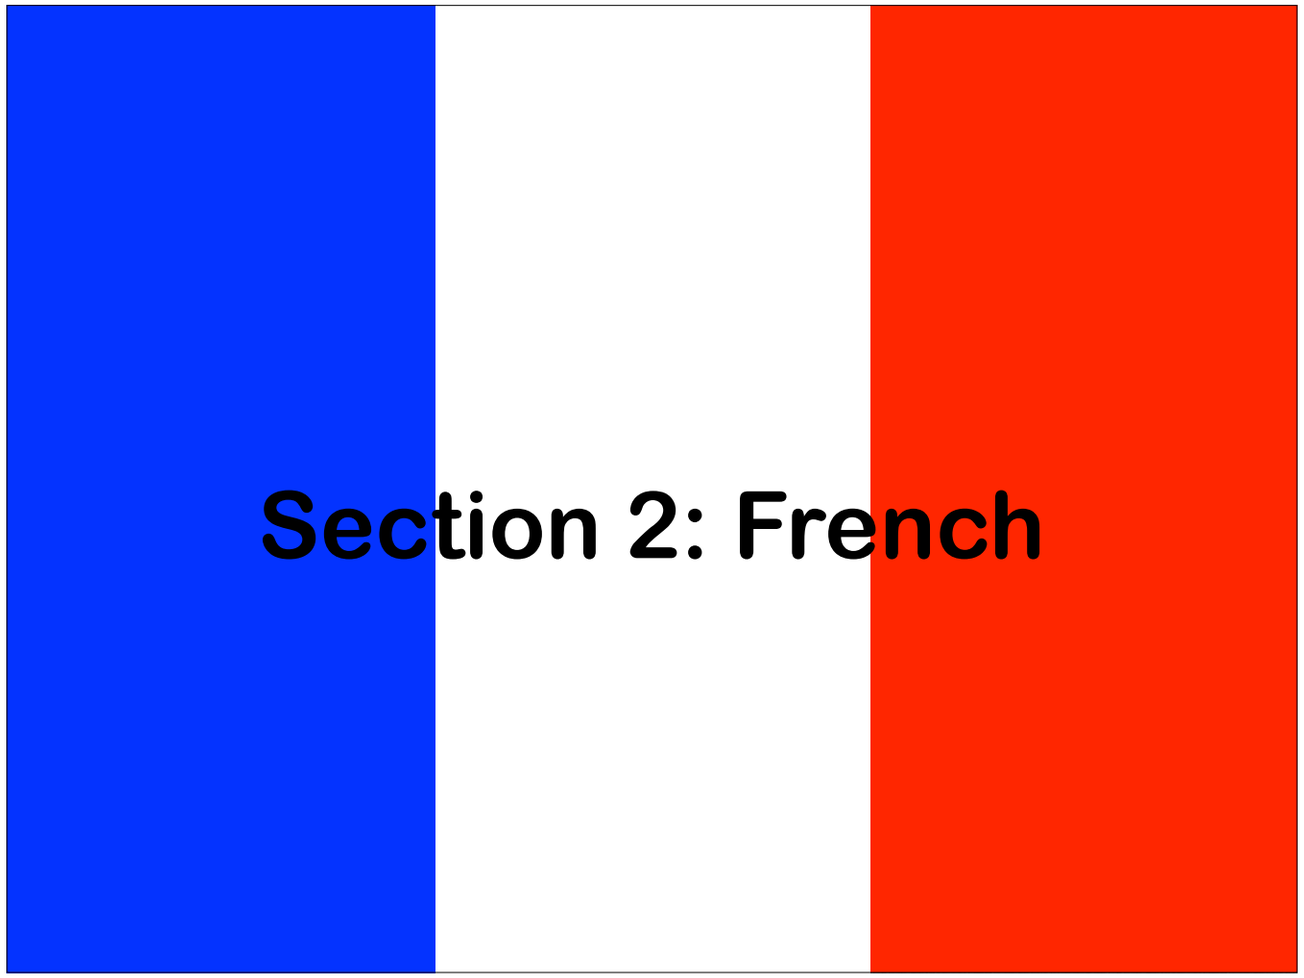

Section 2: French

bête noire

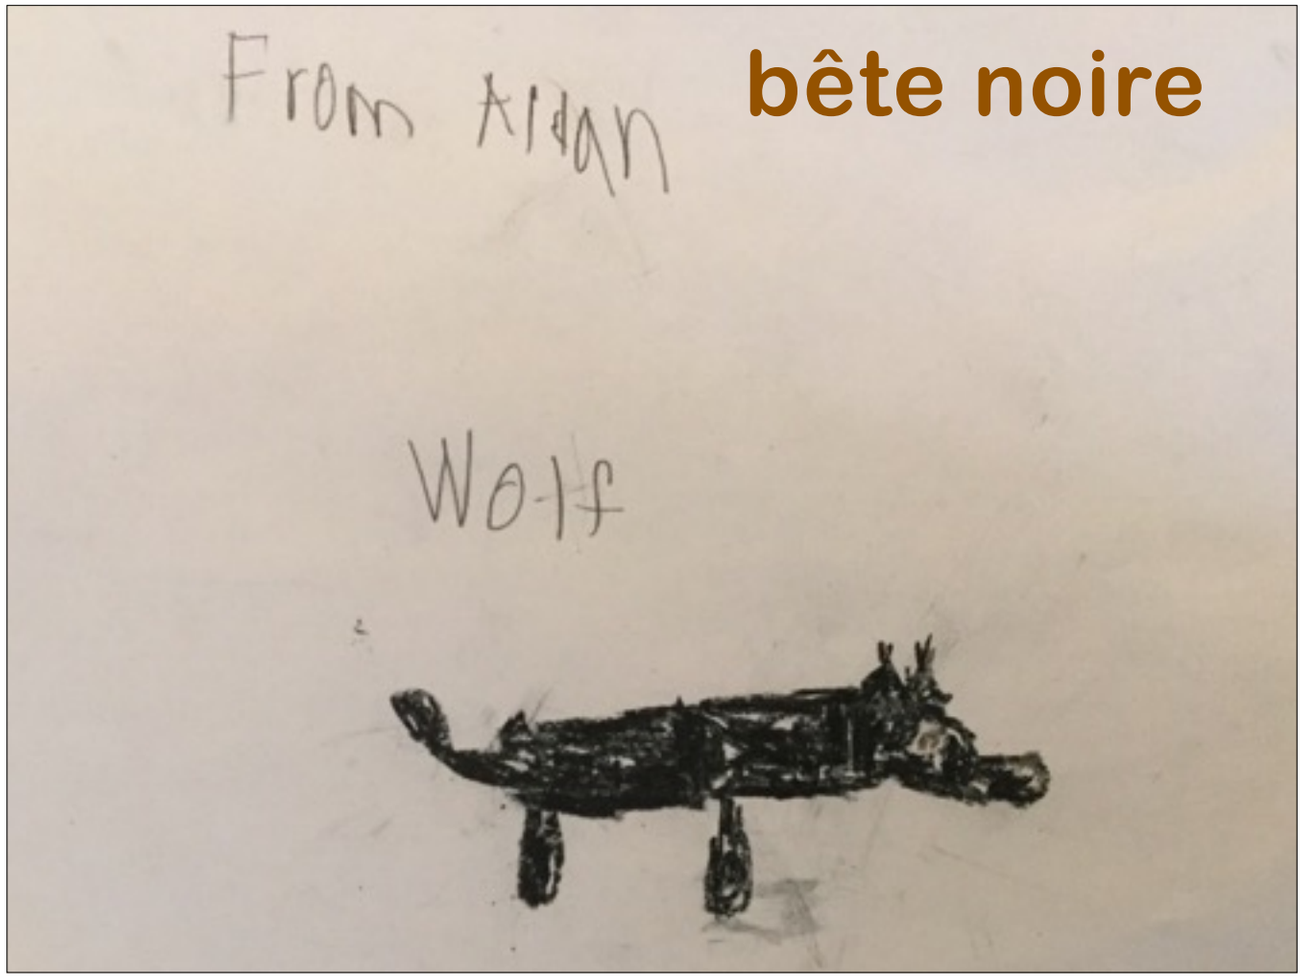

Our French will be limited to one word — bête noire.

(Suggested heckle: That's two words.

Response: That's not an unreasonable point of view, but sub-optimal in my opinion.)

We can say "black beast" perfectly well in English. But it doesn't have the same punch as <<bête noire>>. And I want punch because it really kills me when quants make things worse. We're supposed to be the good guys.

Original drawing from Aidan.

Risk parity

Equalize the variance

contribution from the various groups

In risk parity we have some grouping of assets, and we arrange our portfolio so that an equal amount of variance comes from each group.

(The audience found great humor in “risk parity” immediately following bête noire. That was unintended, but welcome nonetheless.)

The next slide is a picture you’ve seen before, which is it?

(The audience nominated inferno, gas can, and bête noire.

There was a vote — certainly an unamerican election since there were more than two parties.

Bête noire was not very popular. The gas can probably had a slight edge.)

We started out saying “risk” but when we got to the detail, it was variance. In other contexts the detail will be tracking error.

When we say “risk” or “empty” we delude ourselves into thinking we are capturing the whole picture. But we’re not.

Risk parity (part 2)

Big weight in **low** volatility assets

implies low expected return

Risk parity has a return problem.

And there's a second problem …

it’s that damned petrol can again. So now what?

Risk parity (part 2)

Big weight in low volatility assets

suggests low expected return

When I said “imply” it meant that I was believing my model. That’s a bad thing to do.

Whorf-Burns Hypothesis

Our models control our perception of reality

A bit of mischief.

There are two types of quant:

- those who believe that their models overly control their thinking
- those who do not believe that their models overly control their thinking

I claim that the first group is correct — they are overly controlled by their models. Furthermore, I claim that the second group is even more controlled by their models than the first group.

Risk parity (part 2)

Big weight in **low** volatility assets

suggests low expected return

But we were talking about risk parity's return problem.

There is an obvious solution …

just lever up.

Photo from istock.

So here’s the plan:

Put yourself at the designated location, and wait for global warming.

I think this is SBS.

Smartpeople
Being
Stupid

(Others might have a harsher opinion of this levering proposition. Far be it from me to be so cynical.)

Here is Aidan being bête noire.

Variance Targeting

Stabilize portfolio volatility by **selling** high volatility assets when their **volatility** increases

Variance targeting isn’t as well known as risk parity. Here we want to keep the volatility of our portfolio constant through time. Variance targeting and portfolio insurance are siblings.

Variance Targeting

This is Wonderful

Variance Targeting

IF

you are alone.

Photo from istock.

Variance Targeting

Otherwise

If lots of people are doing this, then

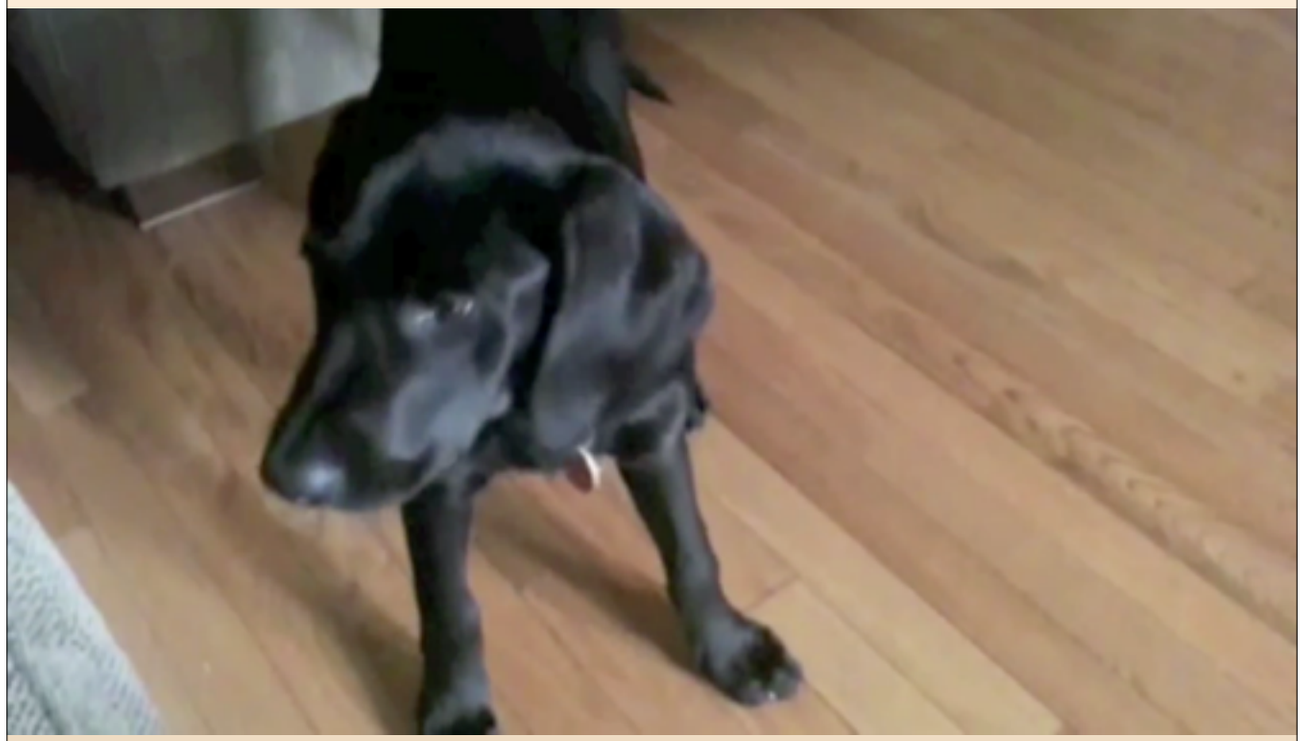

sooner or later a positive feedback loop will be initiated. And how is this feedback loop going to end?

(In the talk there was a 20 second clip of Hunter chasing his tail. Unfortunately the video seems to have disappeared from Youtube, but anyway it doesn't have quite the same effect because it doesn't end at the proper place — with his face smashed into the floor.)

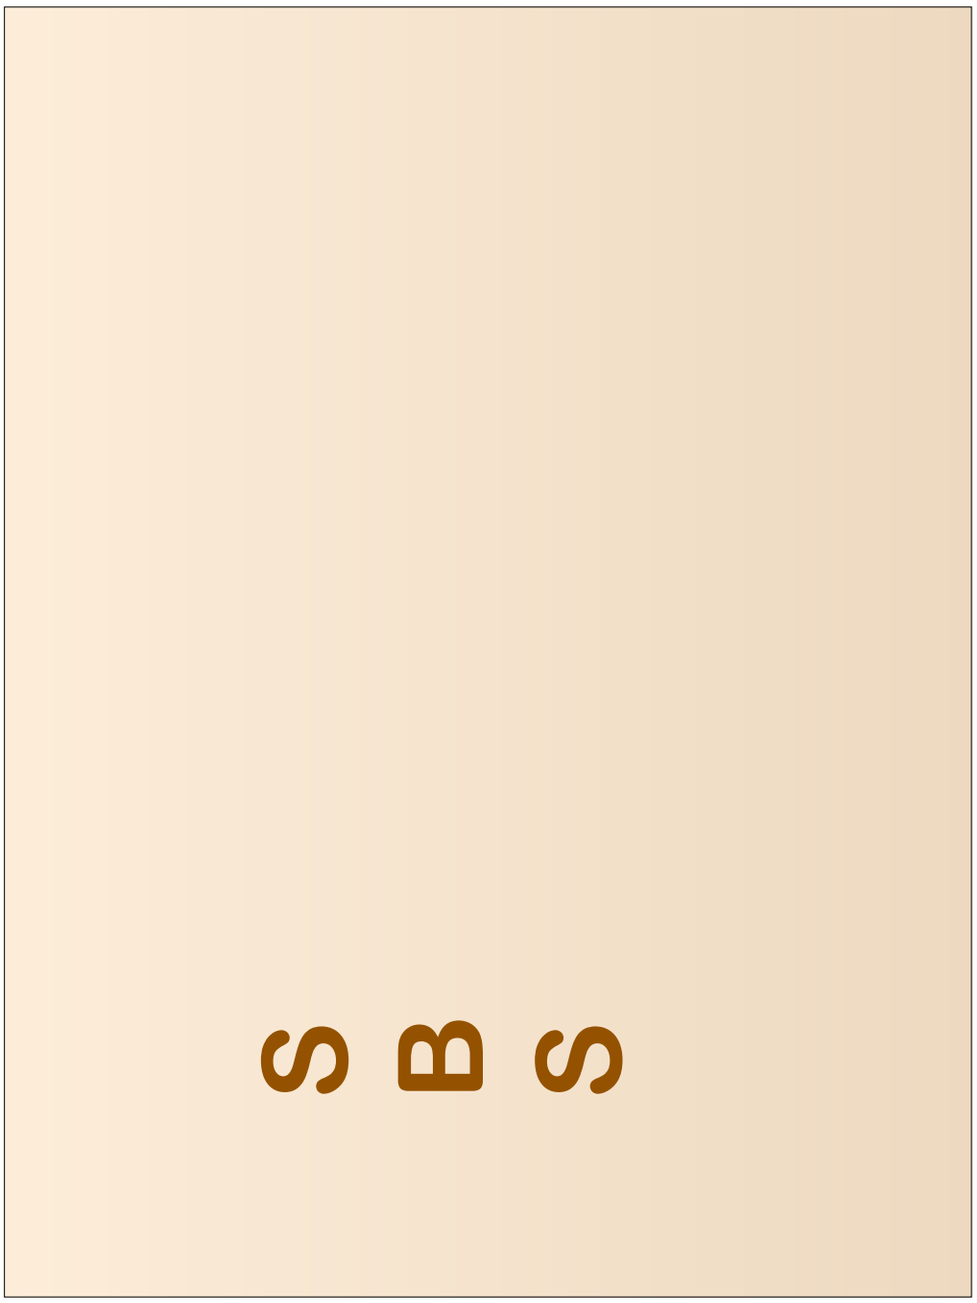
S B S

We are used to walking on solid ground. We are especially used to mathematics being on solid ground — that's the whole point of math.

So quants in particular have a hard time because

Photo from istock.

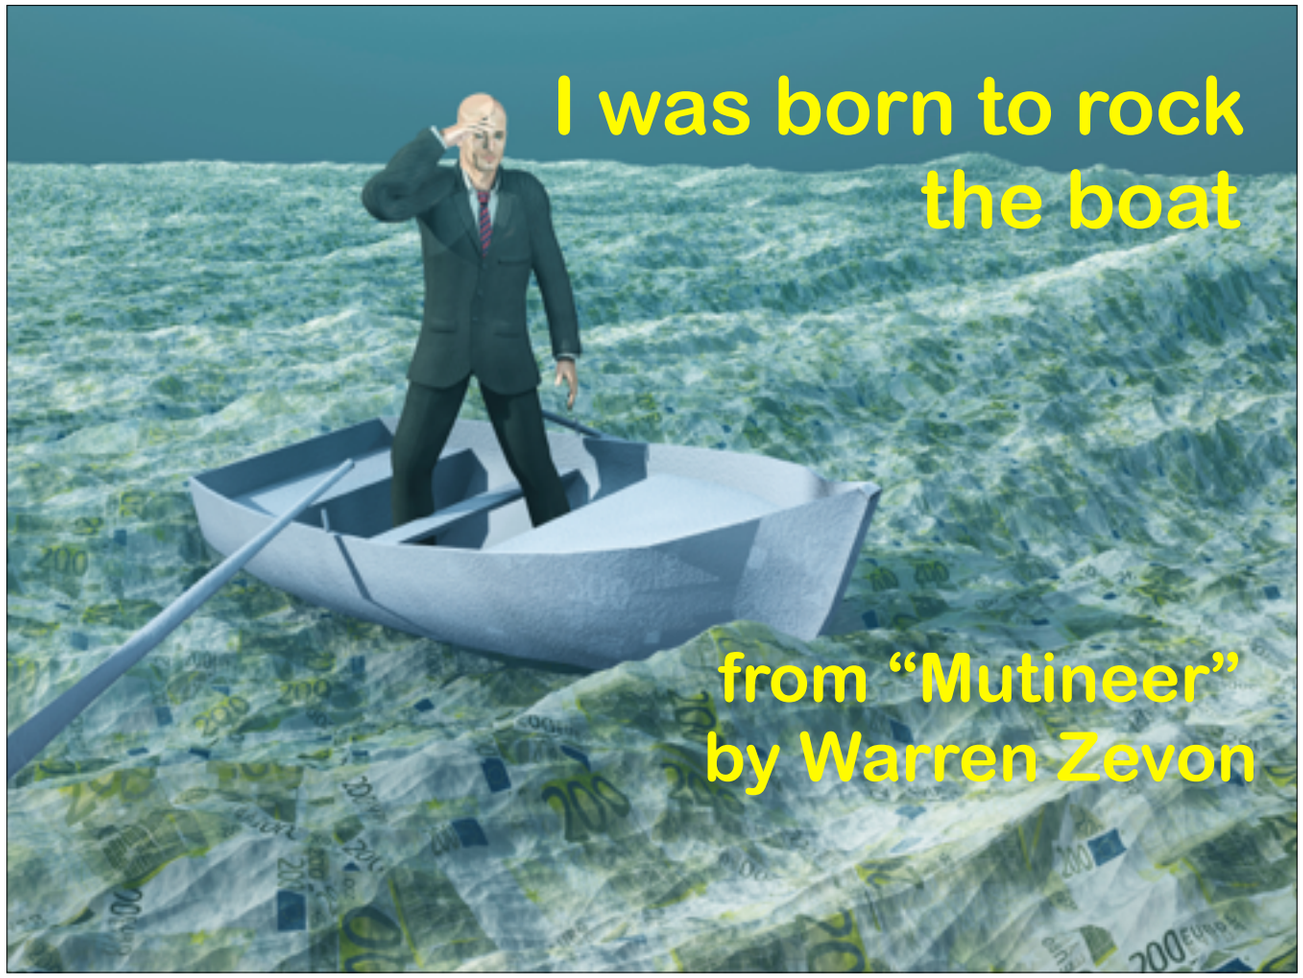

markets are not like that. Markets have a liquidity to them. And the less of the traditional form of liquidity there is, the more of this other liquidity we have.

If we do anything at scale, or cause anything to be done at scale, then we change the market.

Image from istock.

My Great
trading strategy

I have a great trading strategy.

The definition of “great” is that it’s good for me. There is a slight problem though: it’s bad for all of you.

We call that an externality. We know what to do with externalities:

Photo from istock.

we tax them.

To Do

- Find and use a word for the rock-the-boatness of markets
- Learn how to tax market pollution

Here is the to-do list so far.

Now we’re ready to go off in a different direction.

Section 3: SAS

(There was a good bit of laughter when this slide came up.

There was silence when the previous slide was up. I hadn't expected the audience to actually believe that I would do something as ludicrous as talk about SAS — forgetting of course that I do ludicrous things.)

Section 3:

To Do

* Learn how to tax market pollution

How can we learn how to tax trading strategies?

Let's build an

Agent-Based Model

My answer is to build agent-based models.

Traditionally we assume a model for the market. It is a TR-42 garch model, with sprinkles on top. And there we are.

Agent-based models are different: we make assumptions about the participants of the market, we let them go at it, and see what happens.

Let's build an

Agent-Based Model

in R

This is an R conference, so I'm proposing to do this in R.

Dumb idea?

I thought the “in R” part of that sentence was a dumb idea. However, it occurred to me that 99% of my good ideas turn to dust, so how much worse could it be with a different prior?

(I asked how many in the audience thought the “in R” was a dumb idea. I saw one lone hand go up. I accused the room of irrational exuberance.)

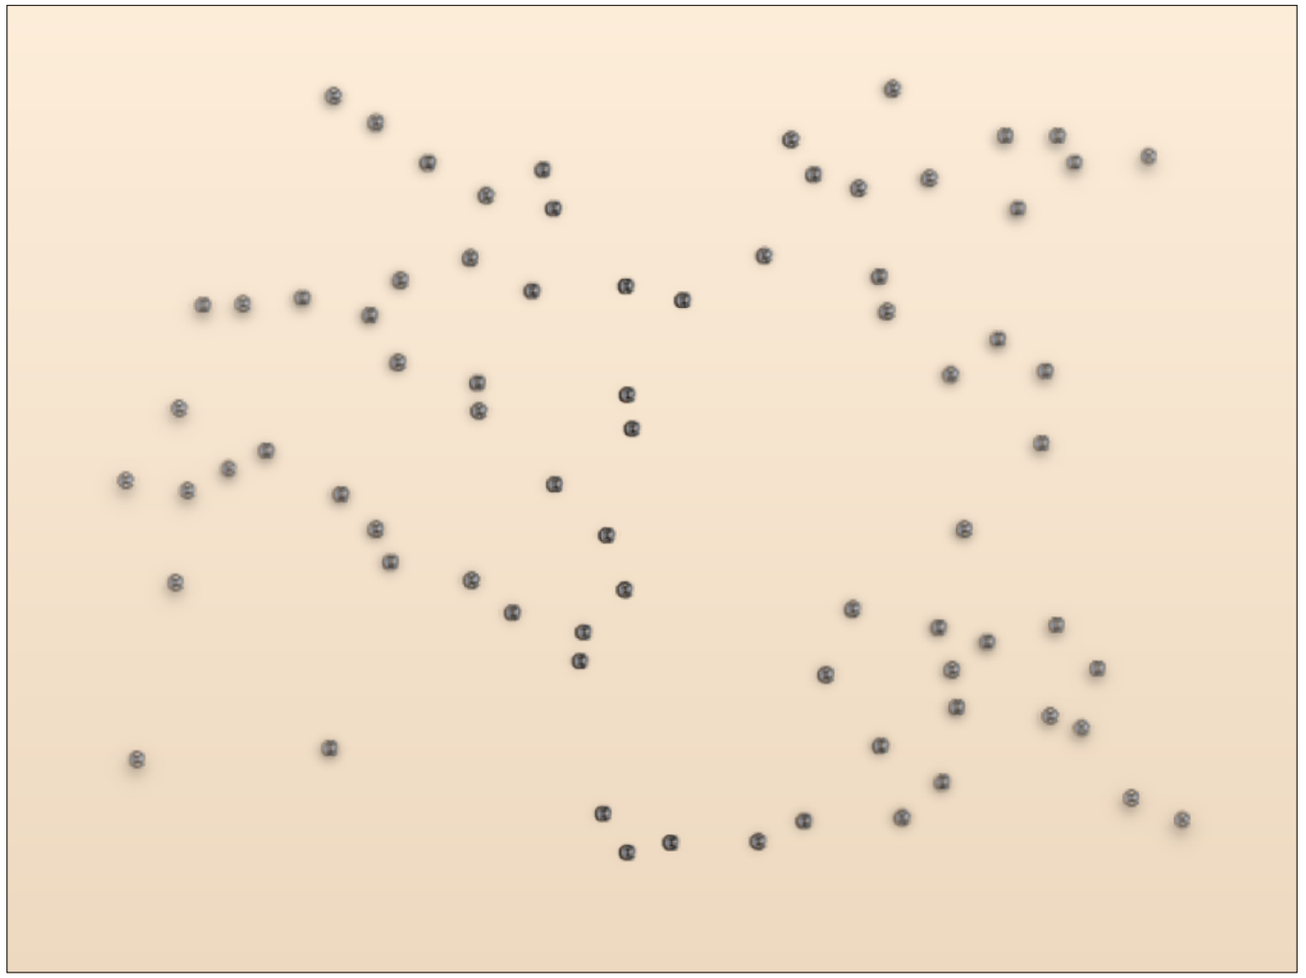

So why did I think it a dumb idea?

We have all these agents scattered around. We need to keep track of what they have and what they intend to do.

This seems like it is playing to R’s weaknesses rather than its strengths.

And I want my function calls to be verbs and not nouns.

Rotate thinking $\pi/2$ radians

But if we reorient our way of thinking …

Rotate thinking π/2 radians

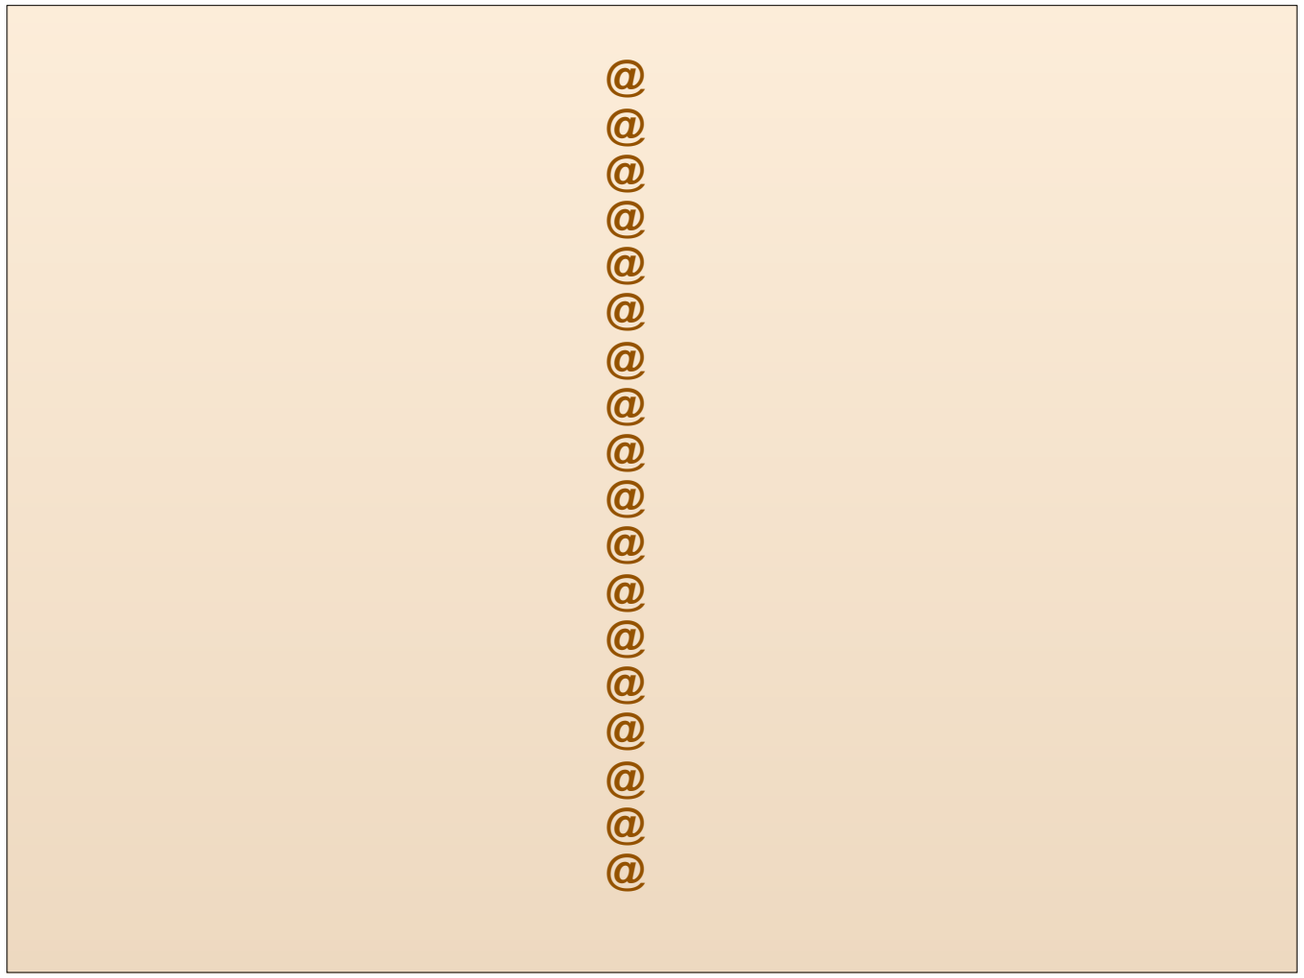

and think of a vector of agents and that vector has a distribution.

Now that is looking like R.

package:marketAgent

Thus was born the marketAgent package.

I'm basically ignorant of agent-based modeling. I'm definitely ignorant of agent-based modeling in R.

And I remained **willfully** ignorant up to the talk. I'm trusting people to give me an education now.

I pointed out to the audience that, given the above, they might consider their strategy for the next few minutes. One strategy would be to pay assiduous attention to what I've done. Another strategy would be to daydream about what they would do.

If I were them, I’d pick strategy number two.

Photo from istock.

Agent components

- cash
- shares
- buyAt
- sellAt
- randomSeed
- call

An agent object contains a vector of cash held by each agent. A vector of shares held. And vectors of prices at which the agents are willing to buy and sell. For example, the first agent might have 2000 in cash and 340 shares; that agent is willing to buy at 89.3717 and willing to sell at 93.8234. Other agents will have different holdings and different value opinions.

The object also has some housekeeping components. Saving the random seed allows us to reproduce this object at will.

Driving function

```
doRun <- function(agents, <<more>>)
{
  <<stuff>>
  for(i in 1:times) {
    thisAuction <- auction(agents, <<more>>)
    agents <- backoffice(agents, thisAuction, <<more>>)
    history[i,] <- thisAuction
    ledger[i,] <- accounts(agents)
    <<other stuff>>
  }
  ans <- list(<<yet other stuff>>)
  class(ans) <- “agentRun”
  ans
}
```

The central function is ‘doRun’. Its primary function is to loop over time points. At each time an auction is held. Then the trades are settled.

Importantly we save the history. Actual market participants pay attention to history — quants certainly do. We want our agents to have that same ability to make their decisions based on history.

The result of ‘doRun’ has an S3 class.

Agents are either all in cash or essentially all in shares once they have traded. We could imagine having two types of cash — cash willing to be traded (as here), and cash held in reserve that can be moved to the first type when conditions are right.

Auction at each time

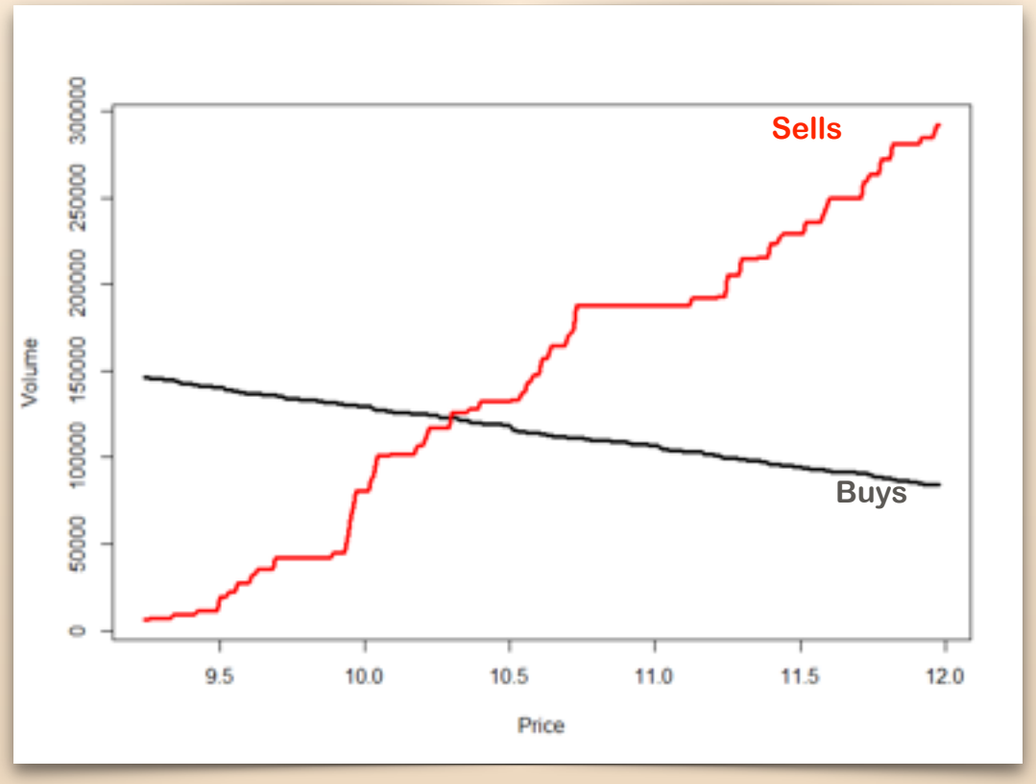

In the auction we want to maximize the volume. We do that by picking the price where the two lines cross.

(I think it was somewhere around here that I heckled the audience for not heckling.)

```
5
6  set.seed(246)
7  a1 <- makeAgents("a")
8  r1 <- doRun(a1)
9  plot(r1)
10
```

So let’s put it into action. We create some agents, do a run, and plot the results.

And what do we get?

```
5
6  set.seed(246)
7  a1 <- makeAgents("a")
8  r1 <- doRun(a1)
9  plot(r1)
10
```

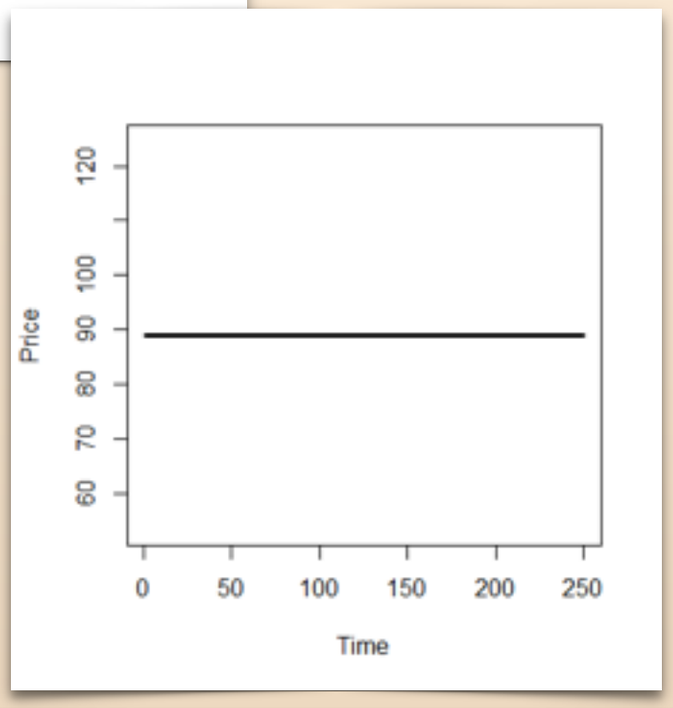

Tada.

I was ecstatic when I got this plot. (The audience probably thought I was joking with that sentence, but it's true: Hey, it works!)

But it does seem like we're missing something.

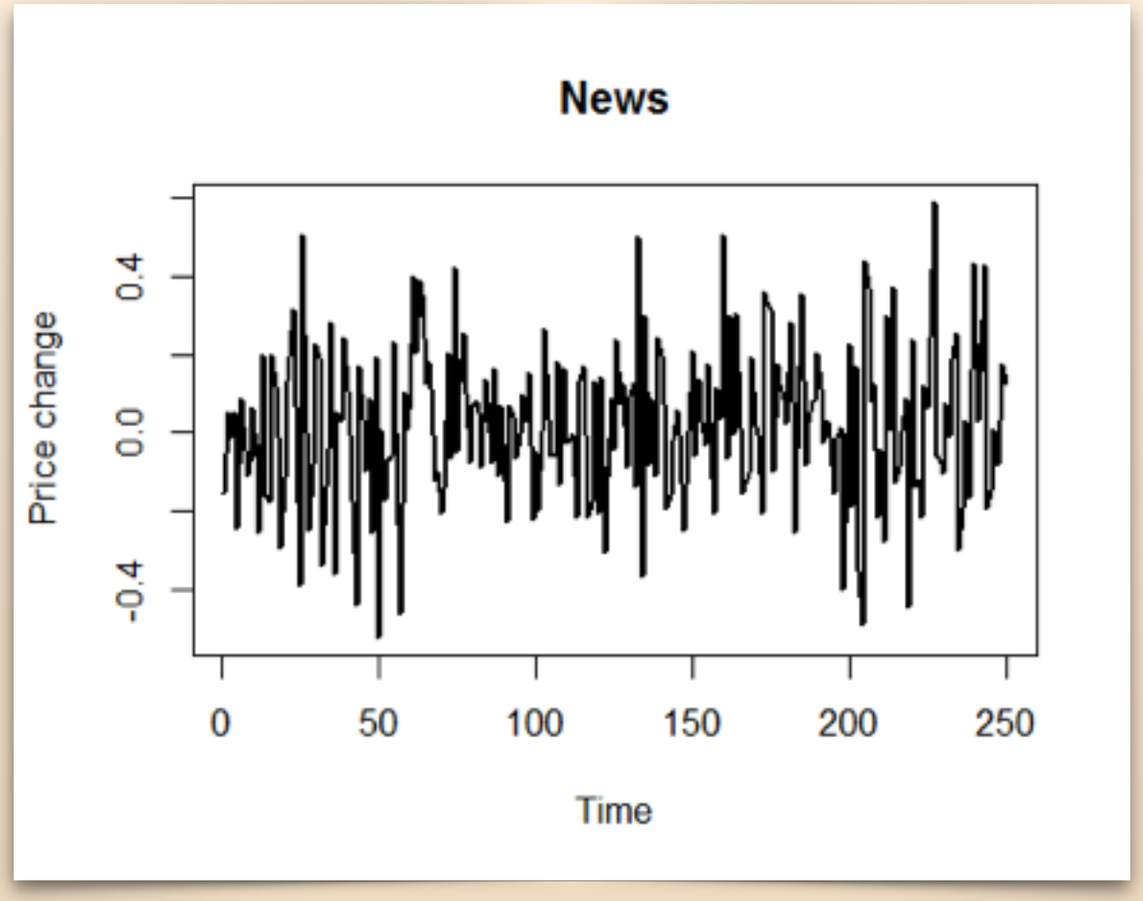

The news says how much the price should change at each time.

```
14
15 set.seed(246)
16 a1 <- makeAgents("a")
17 newsVec1 <- rnorm(250, sd=.2)
18 r1n1 <- doRun(a1, newsSpecific(newsVec1))
19 plot(r1n1)
20
```

Let’s tell the agents the news.

Now what do we get?

```
14
15  set.seed(246)
16  a1 <- makeAgents("a")
17  newsVec1 <- rnorm(250, sd=.2)
18  r1n1 <- doRun(a1, newsSpecific(newsVec1))
19  plot(r1n1)
20
```

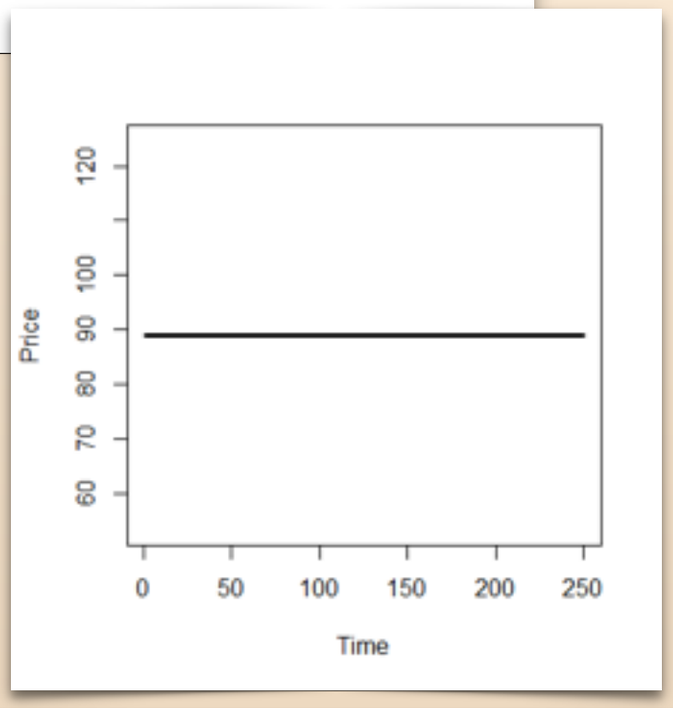

Something scary.

All the agents know precisely how the price should change yet nothing happens at all.

Actually, I'm lying a little. You **might** get a picture like the one here.

```
14
15  set.seed(246)
16  a1 <- makeAgents("a")
17  newsVec1 <- rnorm(250, sd=.2)
18  r1n1 <- doRun(a1, newsSpecific(newsVec1))
19  plot(r1n1)
20
```

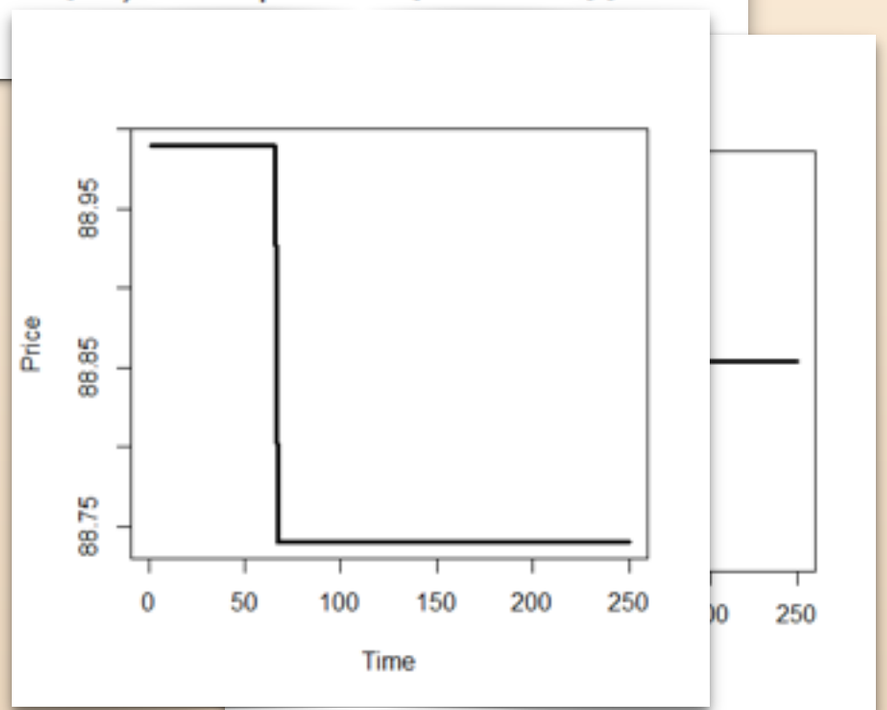

But you are more likely to get a picture like this, where there is one point where the price changes and it is flat everywhere else.

What's happening?

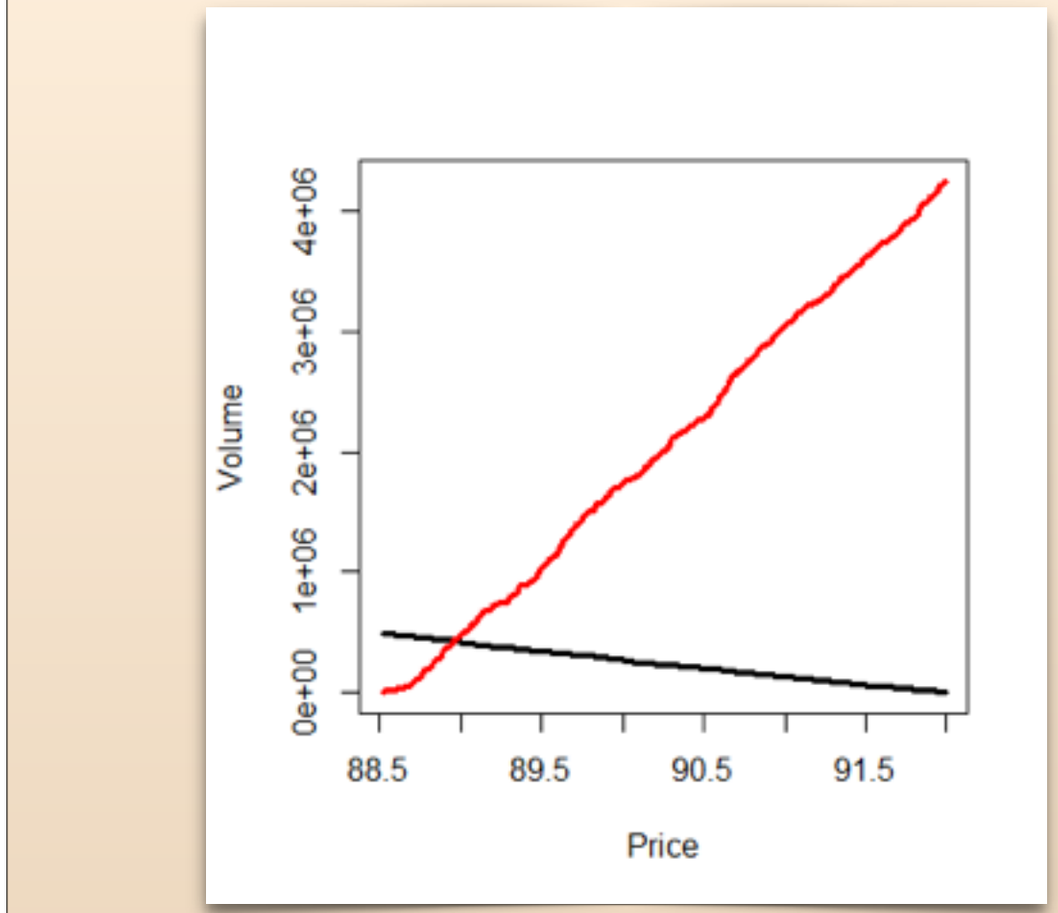

At the first time we'll have a situation like this. But immediately after that, we have

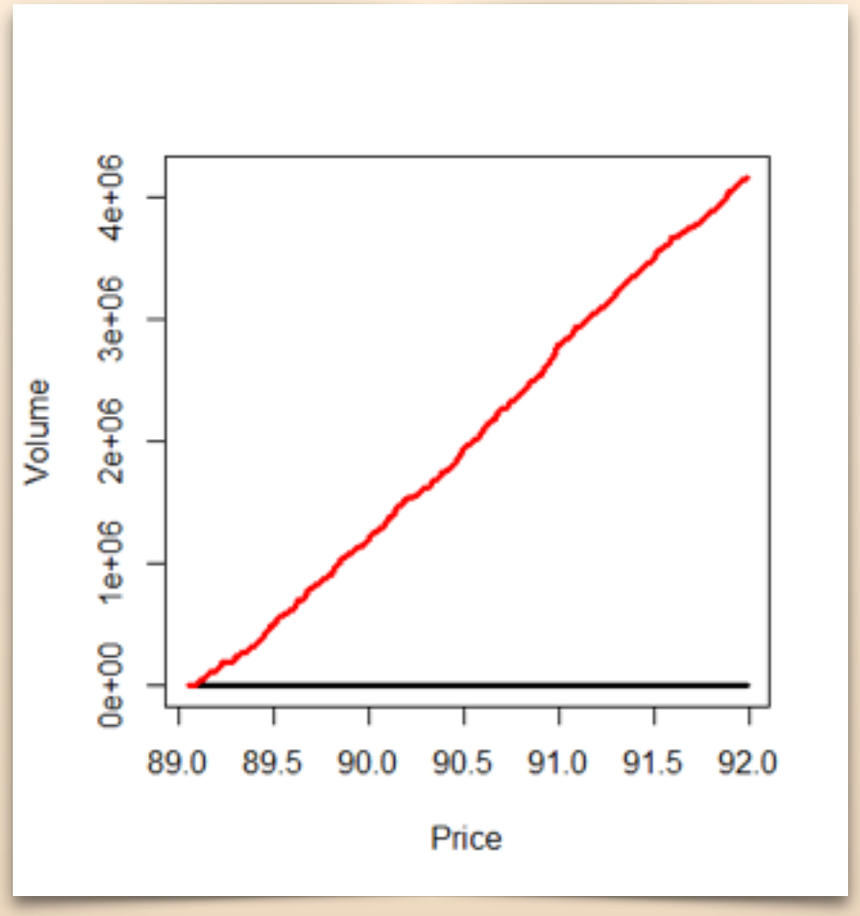

this situation where the lines don’t cross any more.

When we add in the news, we are really just shifting the labels on the x-axis back and forth. The lines don’t move. Well, they almost don’t move.

The trading is done at discrete ticks, but the news and the agent price preferences are continuous. If the news lands just right, then we get a small little crossing of the lines, a small trade, and the price jumps.

So what can we do to make the modeling look reasonable?

(I asked the audience. I was thankful that no one had any ideas because it took me a long time to figure something out.)

```
24
25 marketAgentOptions$traderFraction <- 0.1
26 set.seed(246)
27 al <- makeAgents("a")
28 newsVec1 <- rnorm(250, sd=.2)
29 r1n1pt <- doRun(al, newsSpecific(newsVec1))
30 plot(r1n1pt)
31
```

Note that at each time it is always the same set of agents competing against each other. Maybe that isn't even realistic.

Let's change it so that a random 10% of the agents enter each auction. Now we give the buy and sell lines a chance to cross as the news changes.

```
24
25  marketAgentOptions$traderFraction <- 0.1
26  set.seed(246)
27  a1 <- makeAgents("a")
28  newsVec1 <- rnorm(250, sd=.2)
29  r1n1pt <- doRun(a1, newsSpecific(newsVec1))
30  plot(r1n1pt)
31
```

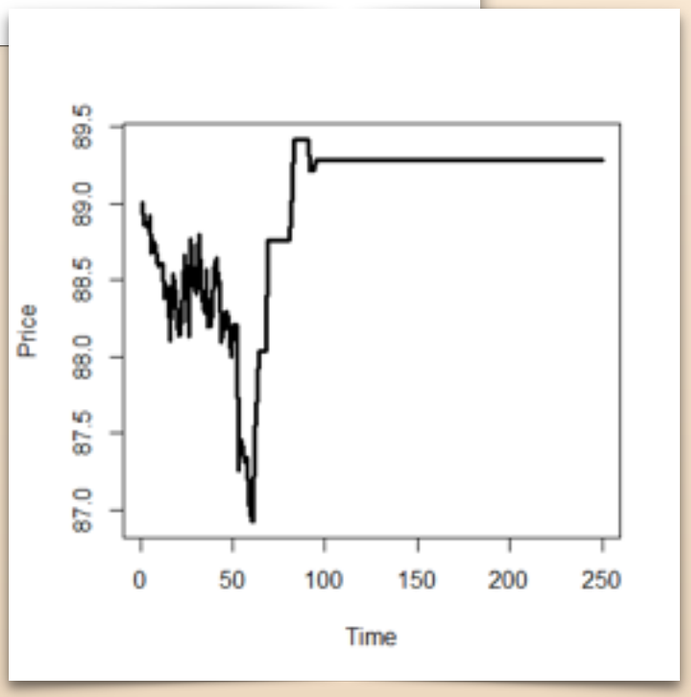

Now at least we get something. At the start it looks rather like market prices. At the end it seems that all the agent subsets have become immune to each other. And there's a transition period.

```
24
25  marketAgentOptions$traderFraction <- 0.1
26  set.seed(246)
27  al <- makeAgents("a")
28  newsVec1 <- rnorm(250, sd=.2)
29  r1n1pt <- doRun(al, newsSpecific(newsVec1))
30  plot(r1n1pt)
31
```

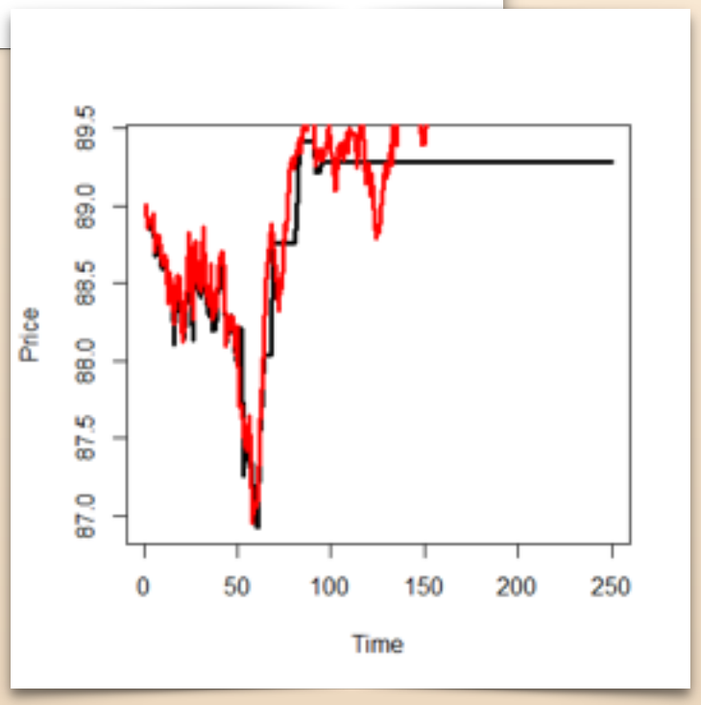

If we add the cumulative news, then we see that indeed the first part is doing the right thing, and even in the transition period it is doing well.

Note that the red line being on top of the black line is merely telling us that the price changes are consistent with the news. We don't have a way of knowing how wrong the initial value is.

```
24
25  marketAgentOptions$traderFraction <- 0.1
26  set.seed(246)
27  al <- makeAgents("a")
28  newsVec1 <- rnorm(250, sd=.2)
29  r1n1pt <- doRun(al, newsSpecific(newsVec1))
30  plot(r1n1pt)
31
```

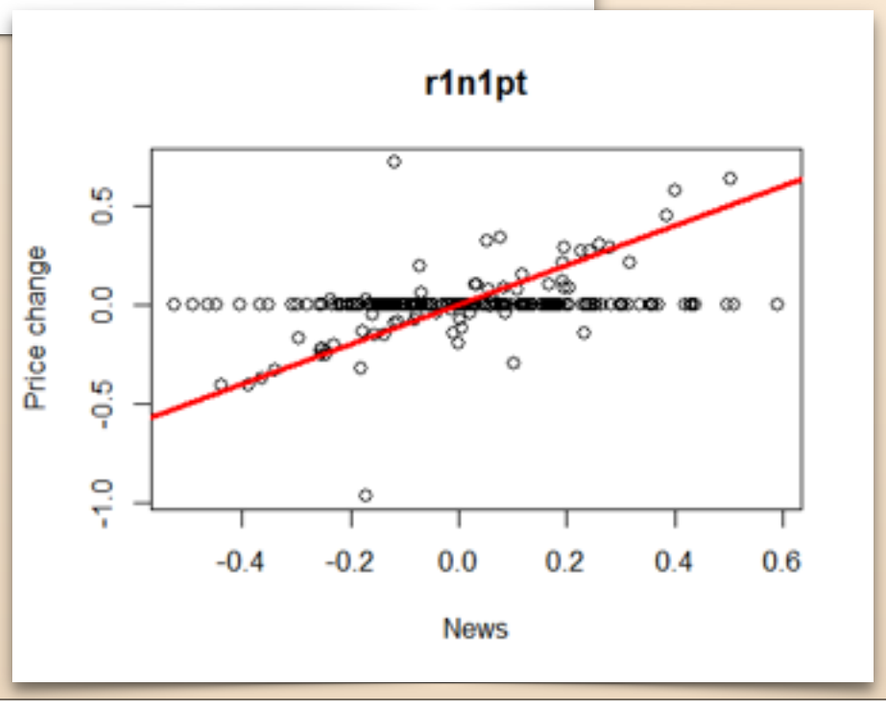

Plotting the news versus the actual market price change shows that there are a few instances where they agree well.

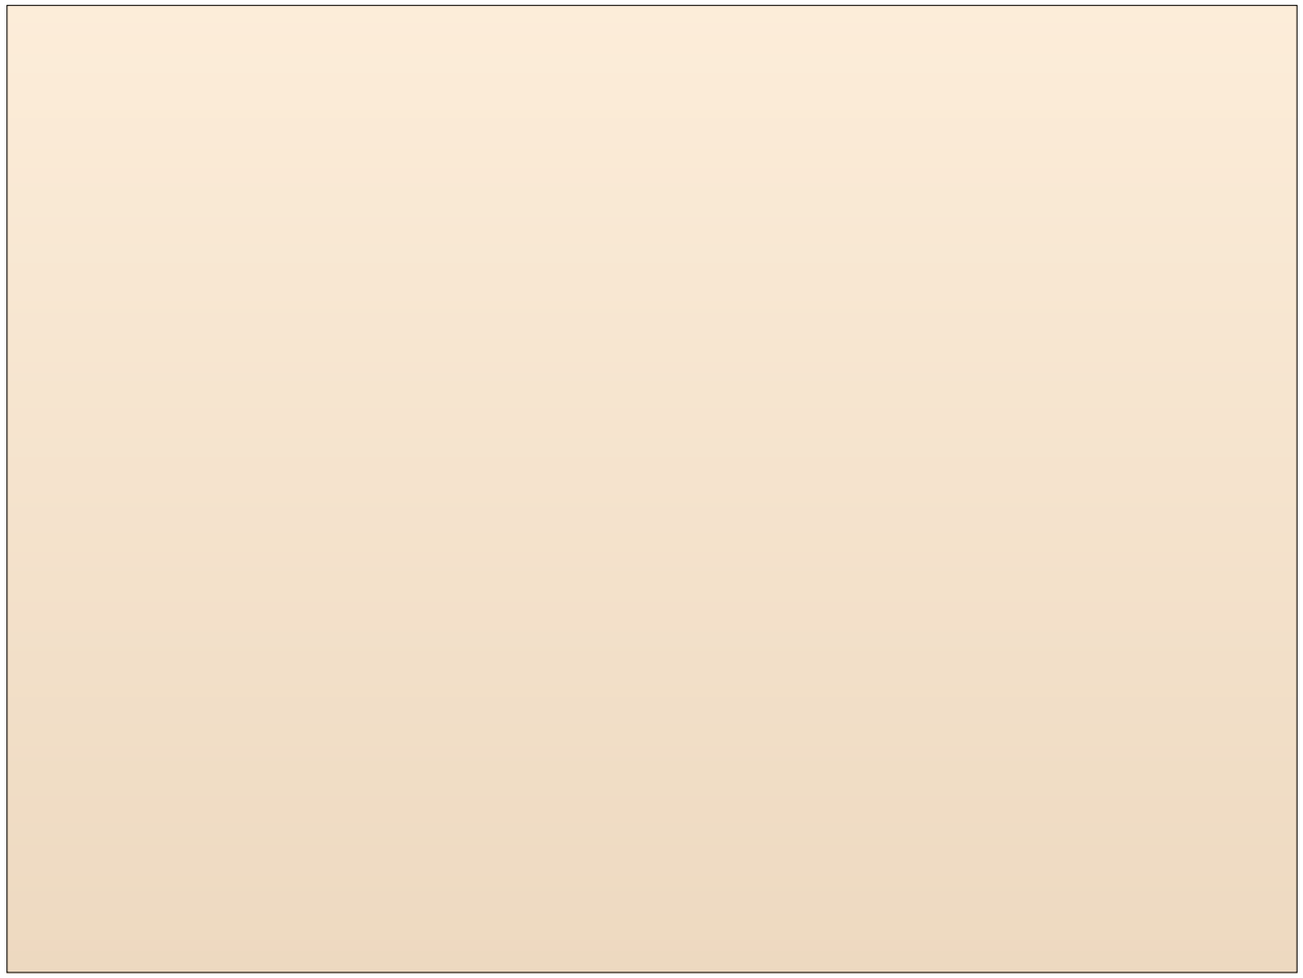

So agents sitting out auctions is a less than a perfect solution.

Well, what if all the agents read the Wall Street Journal, but none of them quite believe it? That is, there is a distribution around the news in the agents beliefs.

```
33
34  marketAgentOptions$traderFraction <- 1
35  marketAgentOptions$privateNewsScale <- 0.05
36  set.seed(246)
37  a1 <- makeAgents("a")
38  newsVec1 <- rnorm(250, sd=.2)
39  r1n1in <- doRun(a1, newsSpecific(newsVec1))
40  plot(r1n1in)
41
```

Let’s go back to everyone always trading, and add in some private news.

```
33
34  marketAgentOptions$traderFraction <- 1
35  marketAgentOptions$privateNewsScale <- 0.05
36  set.seed(246)
37  a1 <- makeAgents("a")
38  newsVec1 <- rnorm(250, sd=.2)
39  r1n1in <- doRun(a1, newsSpecific(newsVec1))
40  plot(r1n1in)
41
```

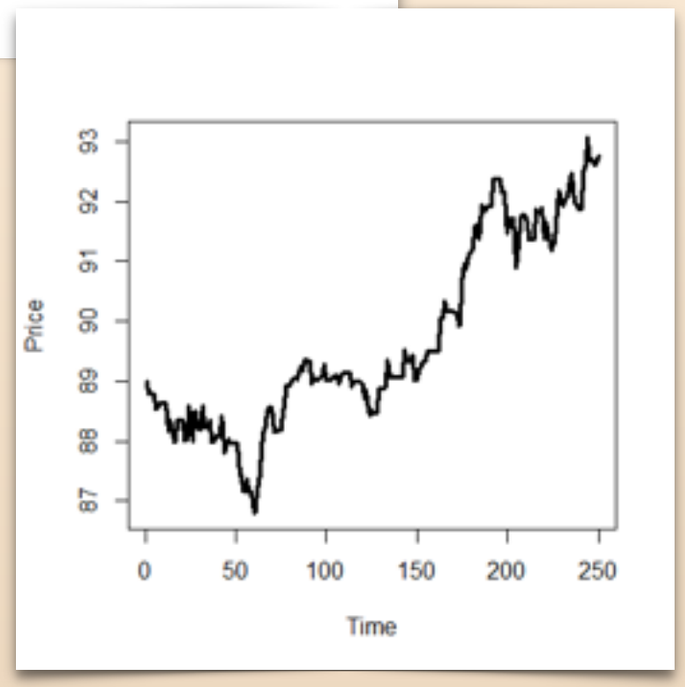

We’re more hopeful with this.

```
33
34  marketAgentOptions$traderFraction <- 1
35  marketAgentOptions$privateNewsScale <- 0.05
36  set.seed(246)
37  a1 <- makeAgents("a")
38  newsVec1 <- rnorm(250, sd=.2)
39  r1n1in <- doRun(a1, newsSpecific(newsVec1))
40  plot(r1n1in)
41
```

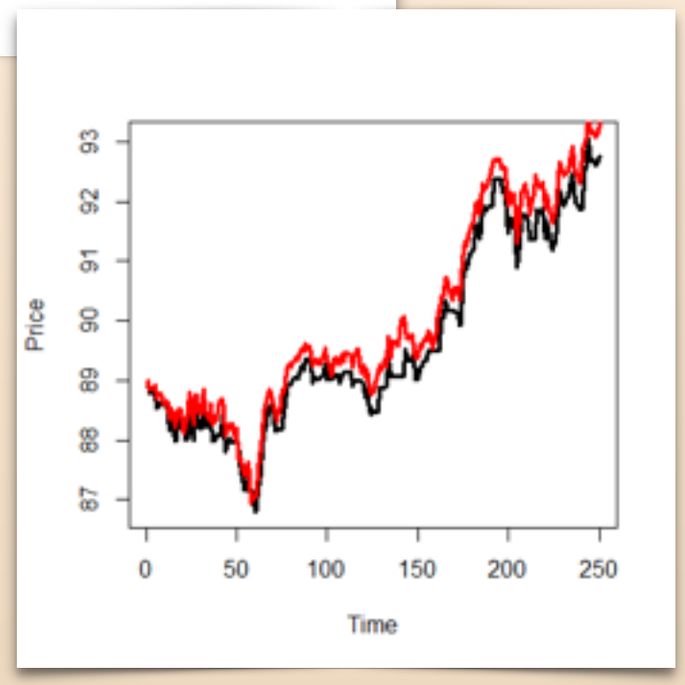

And rightfully so.

I find this weird, wonderful, and a bit scary. We take a machine that is broken, add some noise to it, and it works.

My watch doesn’t work. I know, I’ll smash it with a hammer — now it works.

```
33
34  marketAgentOptions$traderFraction <- 1
35  marketAgentOptions$privateNewsScale <- 0.05
36  set.seed(246)
37  a1 <- makeAgents("a")
38  newsVec1 <- rnorm(250, sd=.2)
39  r1n1in <- doRun(a1, newsSpecific(newsVec1))
40  plot(r1n1in)
41
```

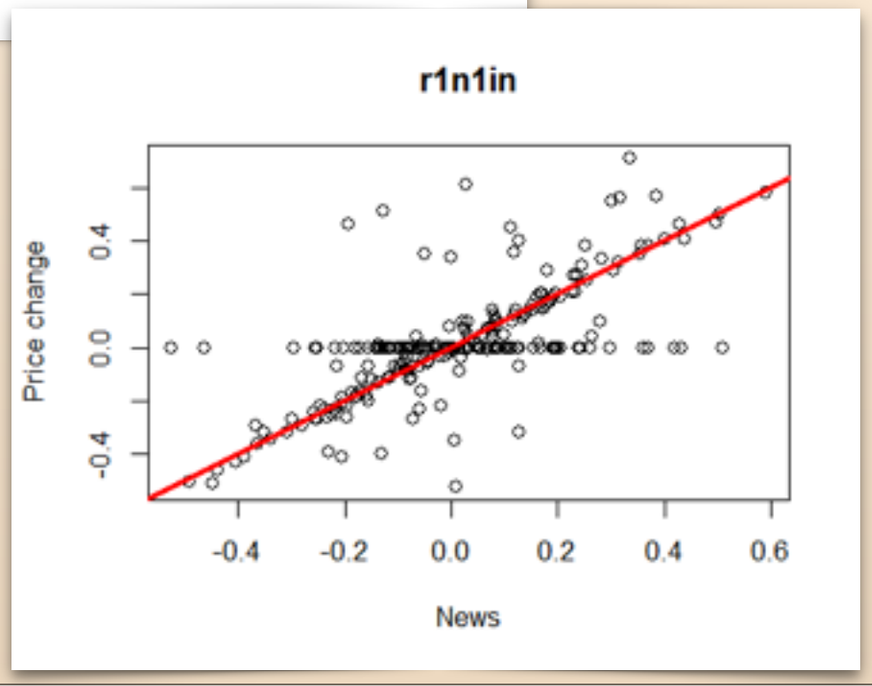

Now there are many more times where news and market change agree well.

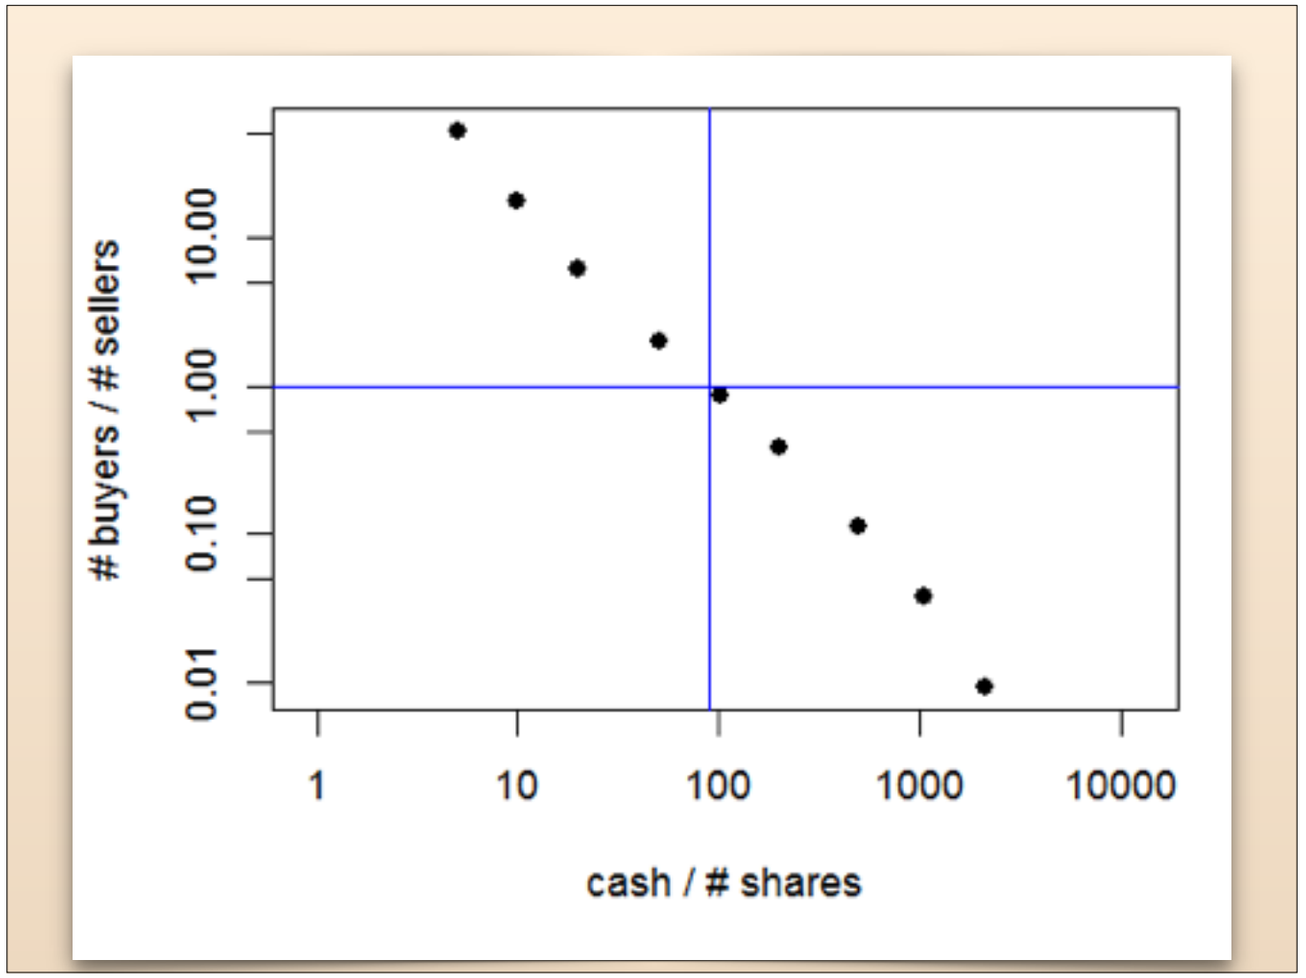

This and the next slide are a bit of a technicality, but it raises its head later.

When we are playing god, we endow agents with three things:

- cash
- shares
- a collective sense of the value

The points in this plot are runs where the sense of value is about 90. The y-axis is median number of buys over median number of sells where the median is taken over the times in the run.

When the cash to shares ratio is about 90, we would expect an equal number of buyers and sellers. That is what we observe.

I expected the slope in this plot to be positive rather than negative. More cash means more buyers, right? Wrong. More cash around means that it takes fewer buyers to suck up all the shares on offer.

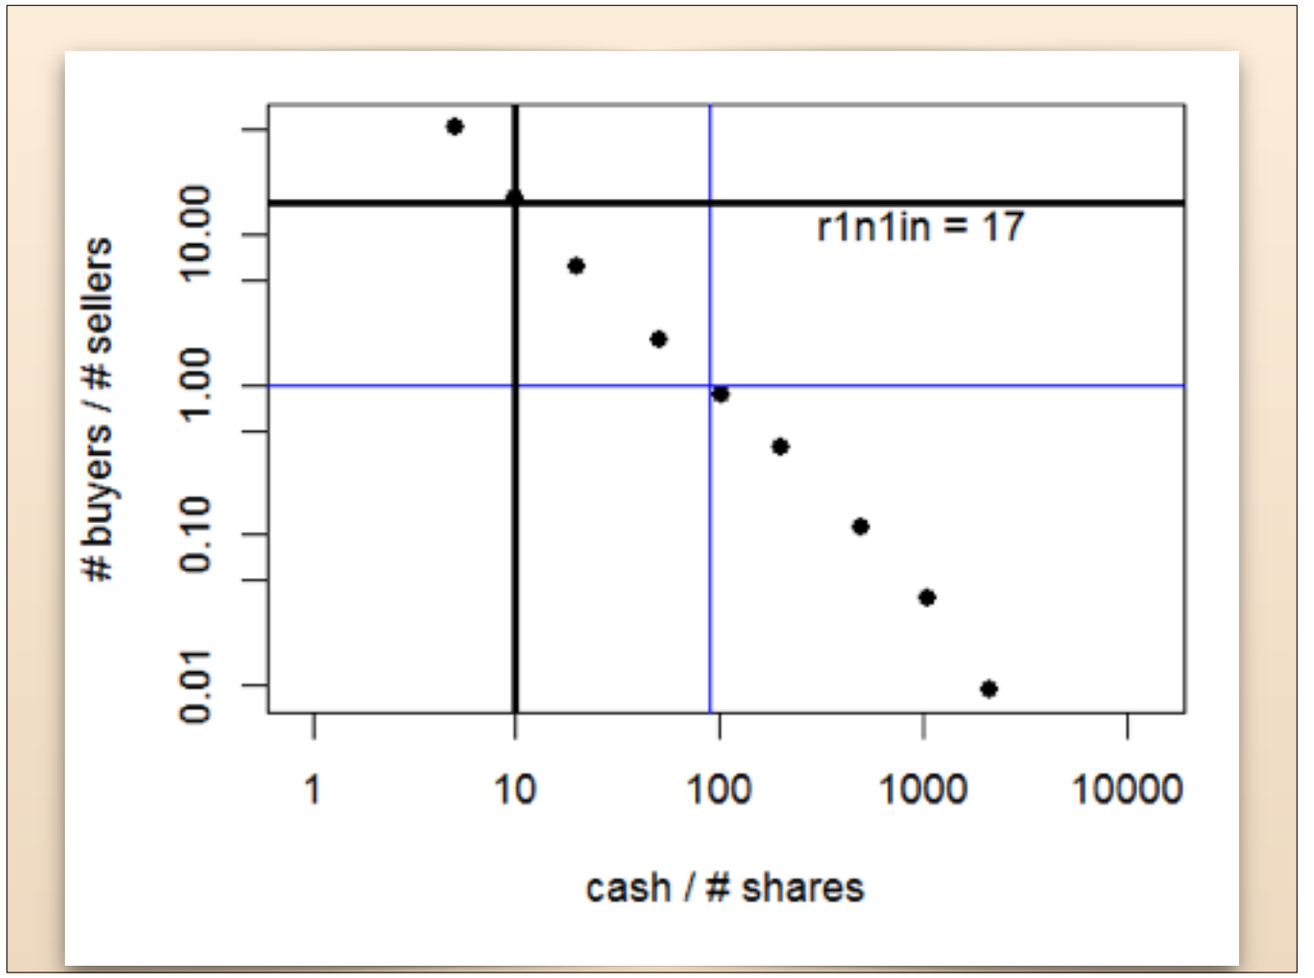

The runs that we have seen and will see have a market price of about 90. But the cash to shares ratio is about 10.

That is because I changed my mind about what the price should be but didn't adjust the cash or shares. That was a stupid thing to do, but it turns out to seemingly have been useful.

The practical thing to note from this is that we ended up with fewer sellers than buyers.

Learn to tax

When I started this I had variance targeting in mind as the strategy to learn to tax. I thought it would be about a century before we would know enough to do the taxing.

I stand by that prediction.

But I realized that there is a strategy that should be easier to learn to tax.

Passive money

If passive induces extra volatility,
then should tax the strategy

That strategy is passive money.

There’s another argument for taxation:

Passive money

If passive induces extra volatility,
then should tax the strategy

Should it be taxed anyway because it is free-riding?

These are completely separate arguments.

Passive money

Market orders

In the agent-based modeling passive is represented by market orders — cash that is willing to trade at any price.

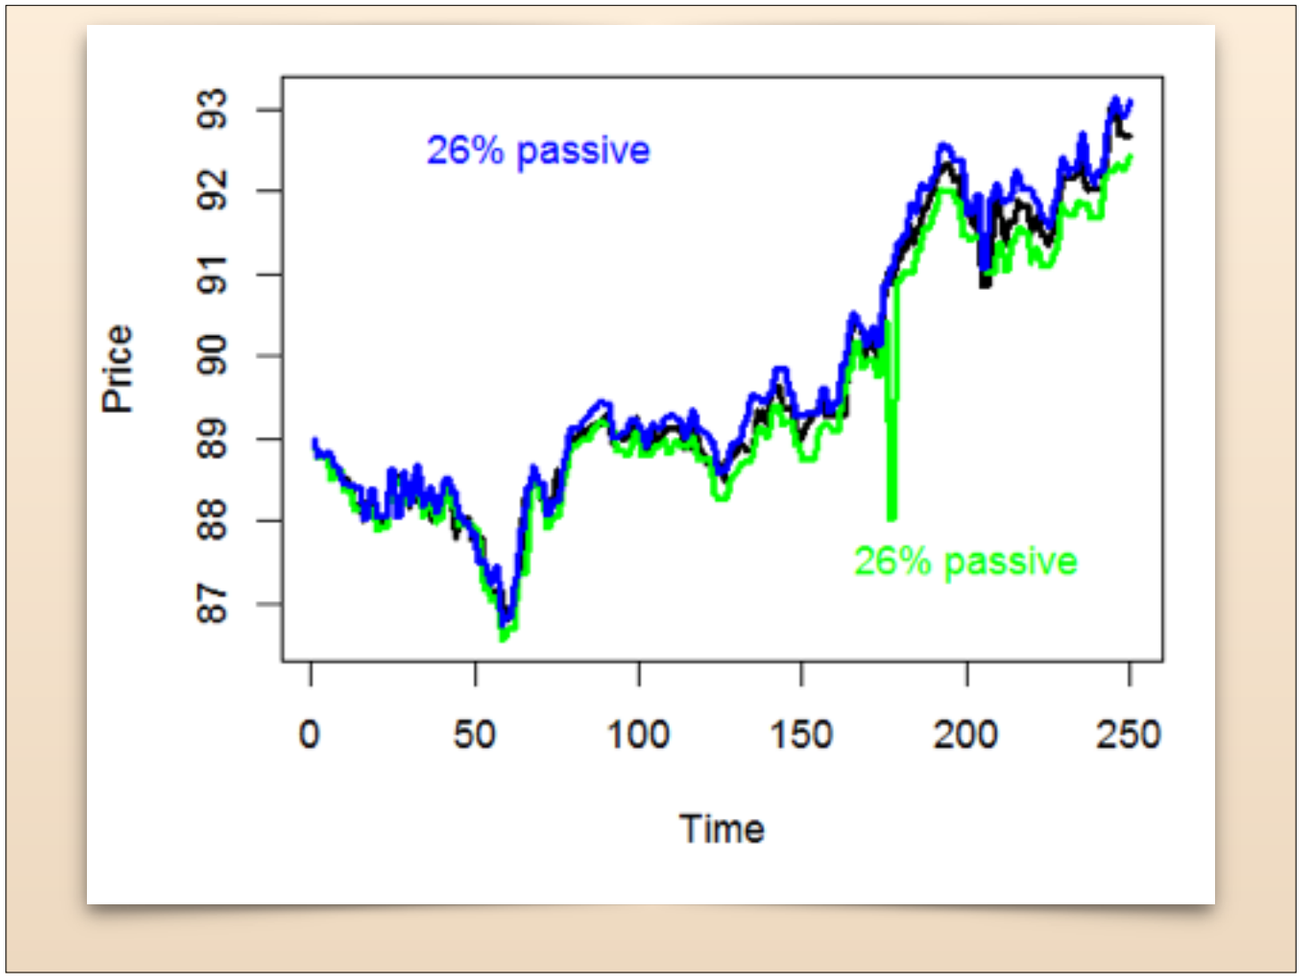

Let’s do an experiment.

The black line is our favorite black line. The blue and green lines have constant pressure to buy and to sell from passive. That pressure results in about a quarter of the volume being from passive.

In general there is very little difference. If this holds true, then the tax for passive should be very very small.

But note that there is a glitch in the green.

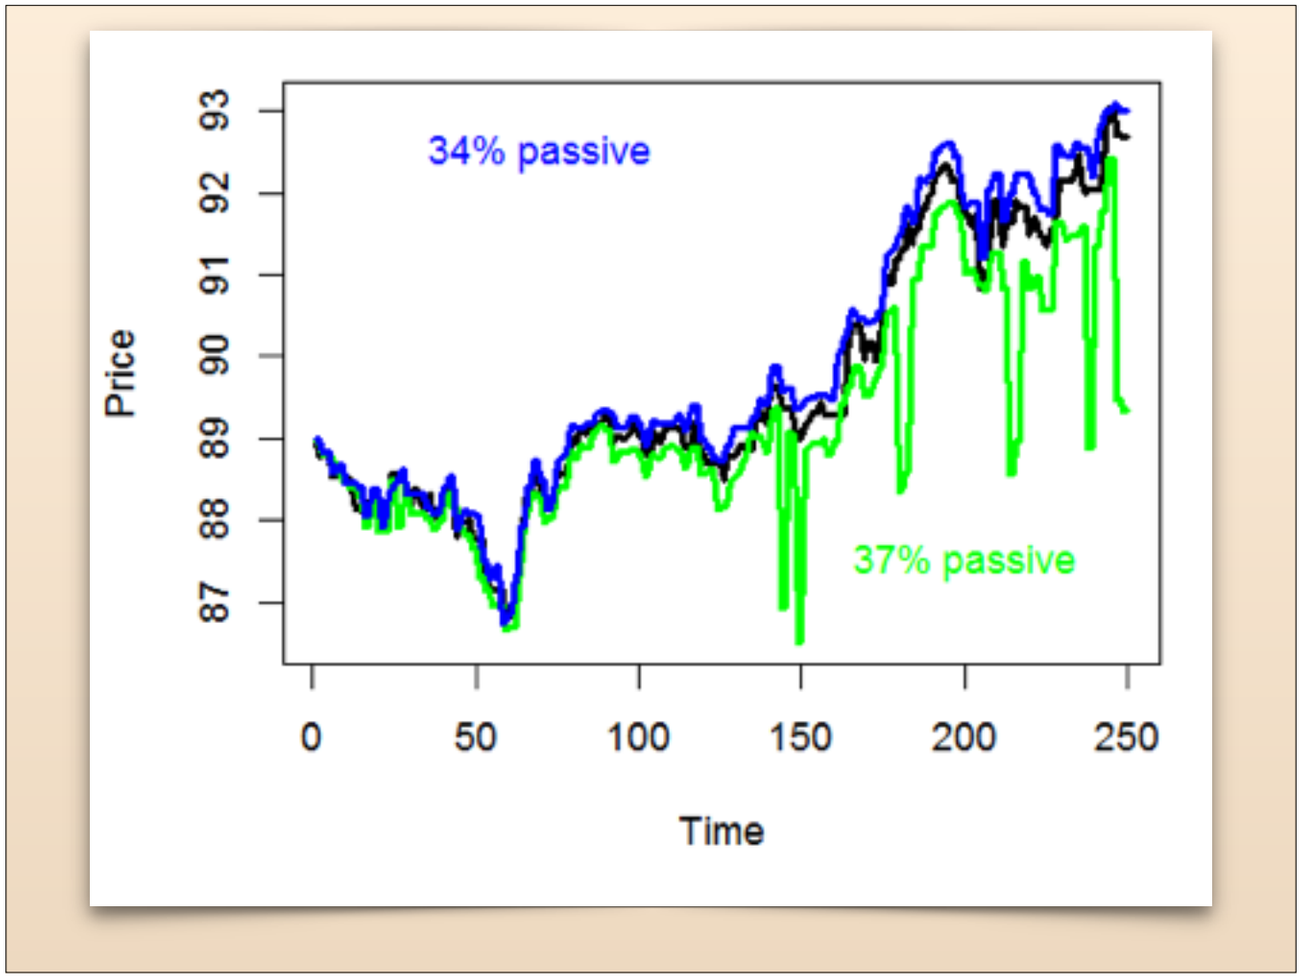

Now we've added more passive pressure. The blue is almost the same, but now we have lots of green glitches. These flash crashes are happening because we are running out of sellers.

During a conversation after dinner on Friday, someone pointed out that if the glitches were pointed up, then we wouldn't think it looked like a real market. (He was talking about some other experiment, he had no idea what I would be talking about.)

I think we are getting a hint of the complexity of the problem even in the easy case.

Markets do **not** react to **news**

I’ve been talking as if markets respond to the news. That is not true.

They respond to our perception of the news. There is no way that I’m going to assume that we are not delusional in our perceptions.

```
readNews.agent_b <-
function(agents, news, history,
         privateNews=getmaOption("privateNews"),
         b_distortion=getmaOption("b_distortion"), ...)
 {
  priv <- match.fun(privateNews)(length(agents$buyAt))
  disnews <- match.fun(b_distortion)(news)
  agents$buyAt <- agents$buyAt + tail(disnews, 1L) + priv
  agents$sellAt <- agents$sellAt + tail(disnews, 1L) + priv
  agents
 }
```

We've already met the A tribe of agents that have an unbiased sense of the news. The B tribe and all of the rest of us have a distorted view of the news.

```
> getmaOption("b_distortion")

function(news) news * 10
```

The B tribe thinks the news is 10 times larger than it actually is. We think of them as an excitable bunch, but perhaps that is not too far off the mark.

So now let's do a run where there are an equal number of A and B tribe members, and with the same news as we've seen before.

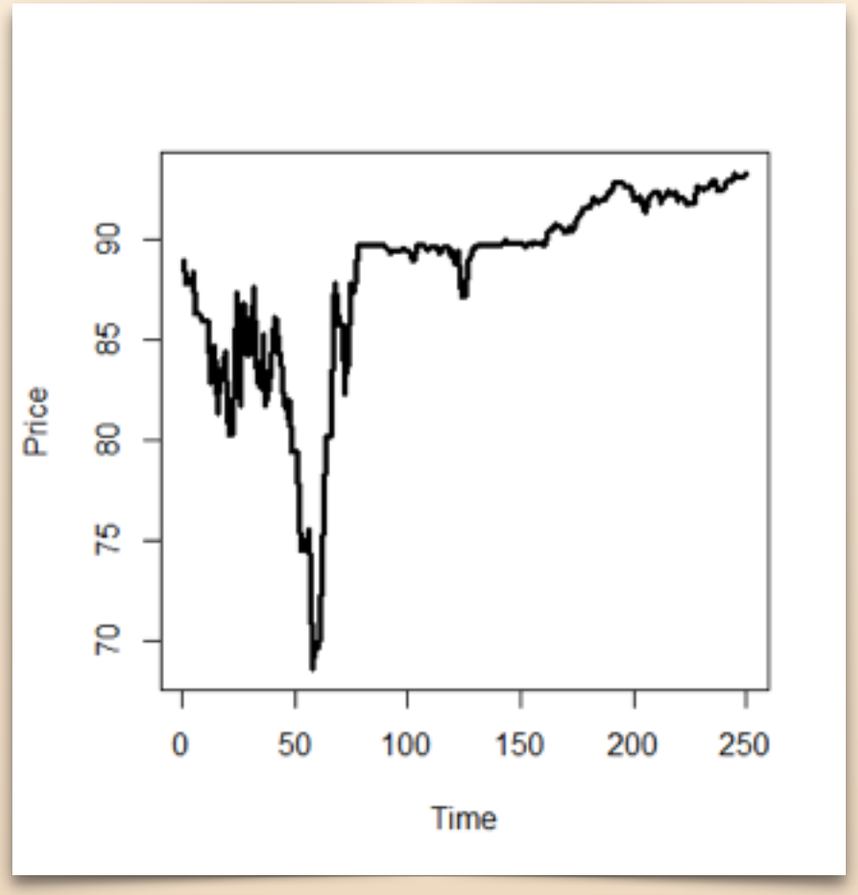

Here’s what it looks like. Any ideas of what is happening?

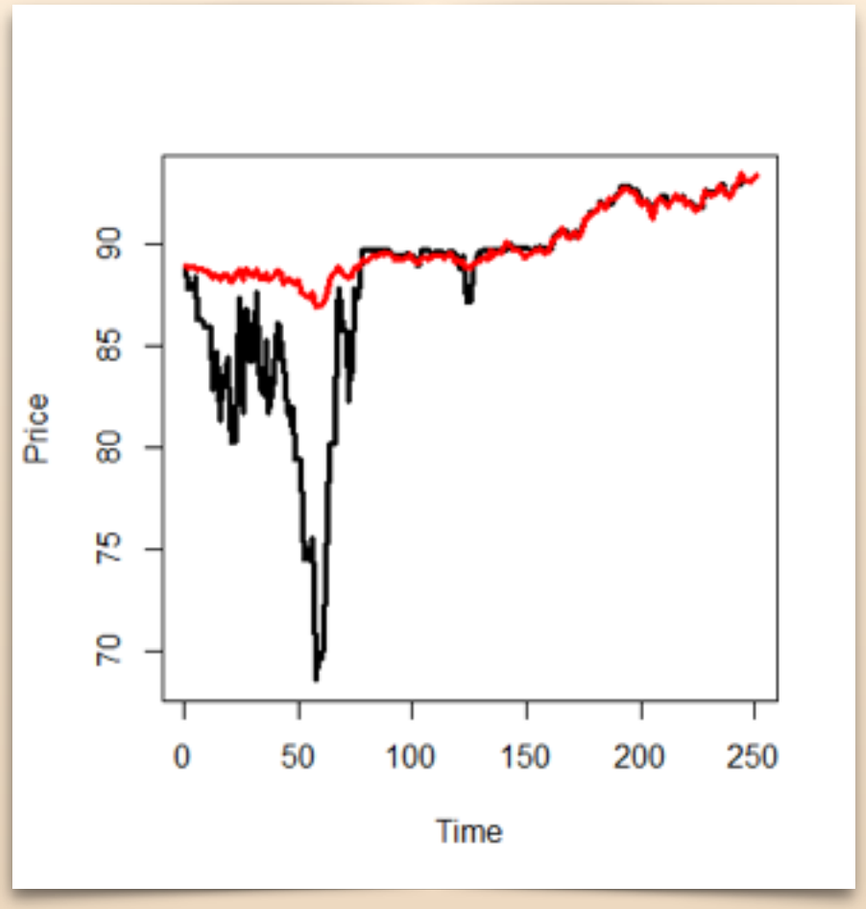

The red line is the actual news. I still didn’t guess what was happening when I got here.

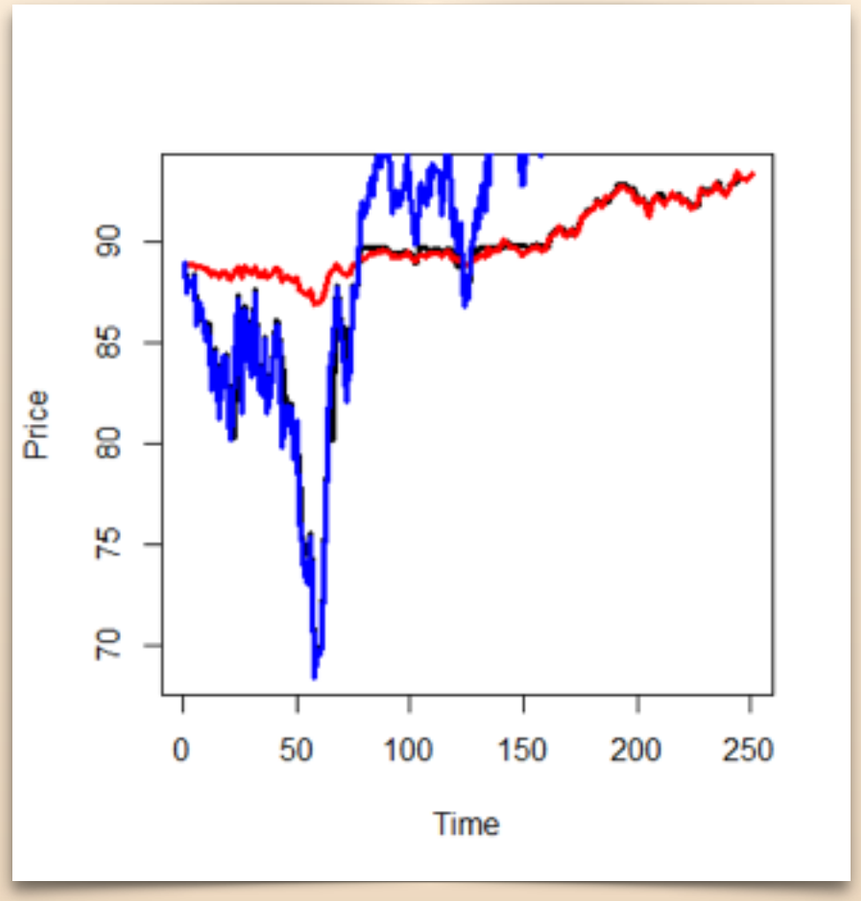

The blue line is the news as seen by the B tribe.

Now we see that the tribes are fighting over the steering wheel. Almost always one or the other has control — hardly ever is there a compromise.

One person pointed out that we are getting the more pessimistic version of the two tribes viewpoints.

Confession Time

Partial fills are troublesome

Share creation/destruction

Currently kludged

I’m hoping that everyone will be shocked to hear me say that the coding in the package is not perfect.

Partial fills are fussy. The phenomenon exists of share/anti-share pairs popping up out of the vacuum, where “anti-share” is more commonly called “cash”. If this is a tiny fraction of the volume at the time, then this shouldn’t be of concern. However, if it is all the volume or a substantial fraction, then it causes headaches. I have a kludgy and probably overly heavy-handed suppression of the problem.

Then of course there are all the bugs that I don’t even know about.

package:marketAgent

- CRAN
- Github
- burns-stat.com

- repos="http://www.portfolioprobe.com/R"

The source file of the package is available via a blog post on the Burns Statistics website. Eventually it will be in the Portfolio Probe repository as well.

It is not on CRAN or Github or any of its coconspirators. It isn't that I don't want it in those places. It's that when I simulate the future, very few of the paths include me spending substantial time on this project.

I would love for someone or someones to take this up and move it from being just a little toy.

The x's and checkmarks are from all-free-download.com

Let’s go back towards reality.

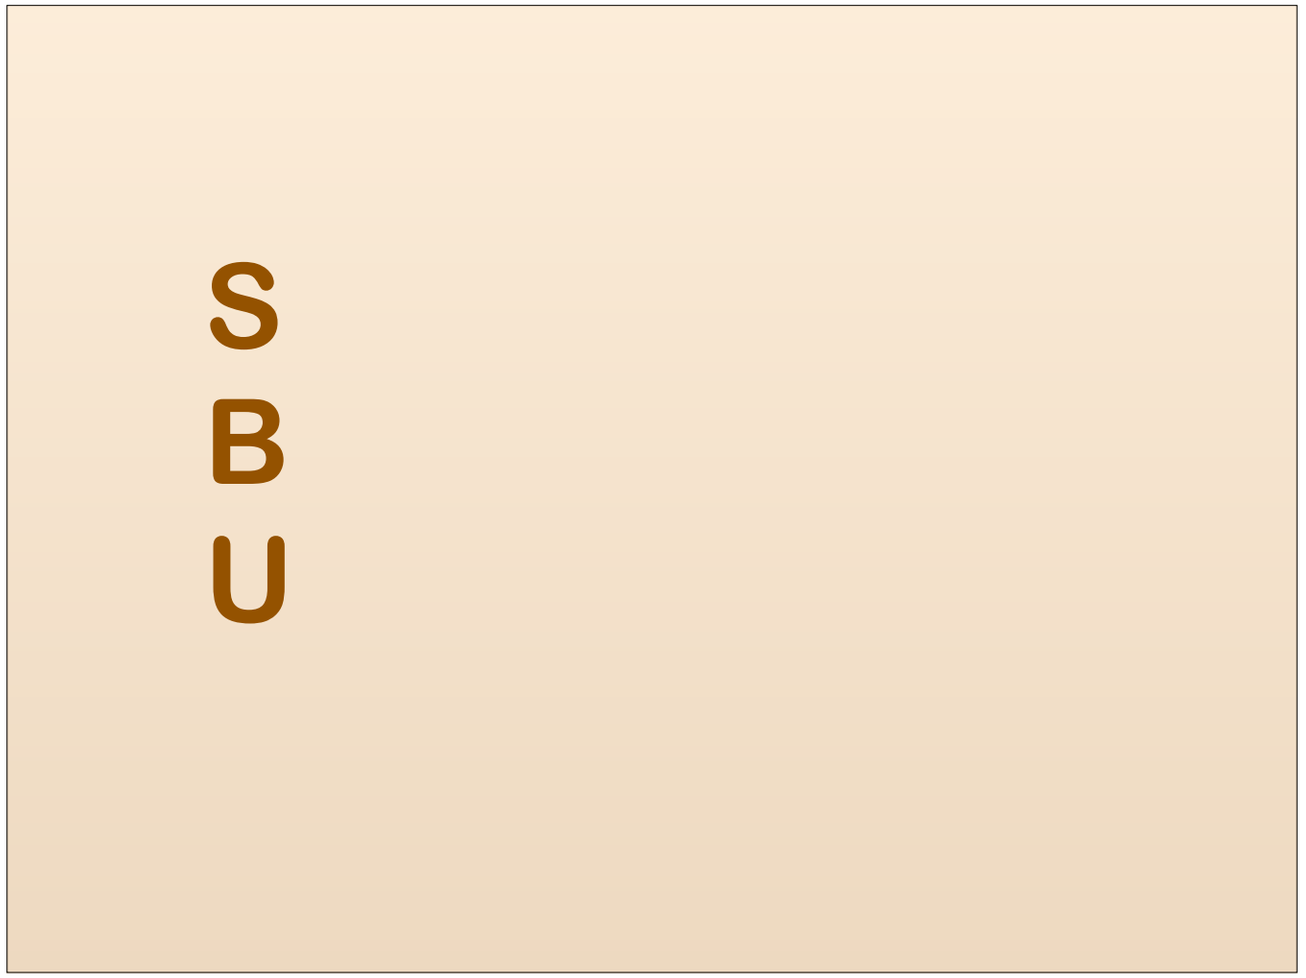

My wish for all of us is that we be SBU.

Smartpeople
Being
Useful

But I leave the final word to one of my favorite quant linguists:

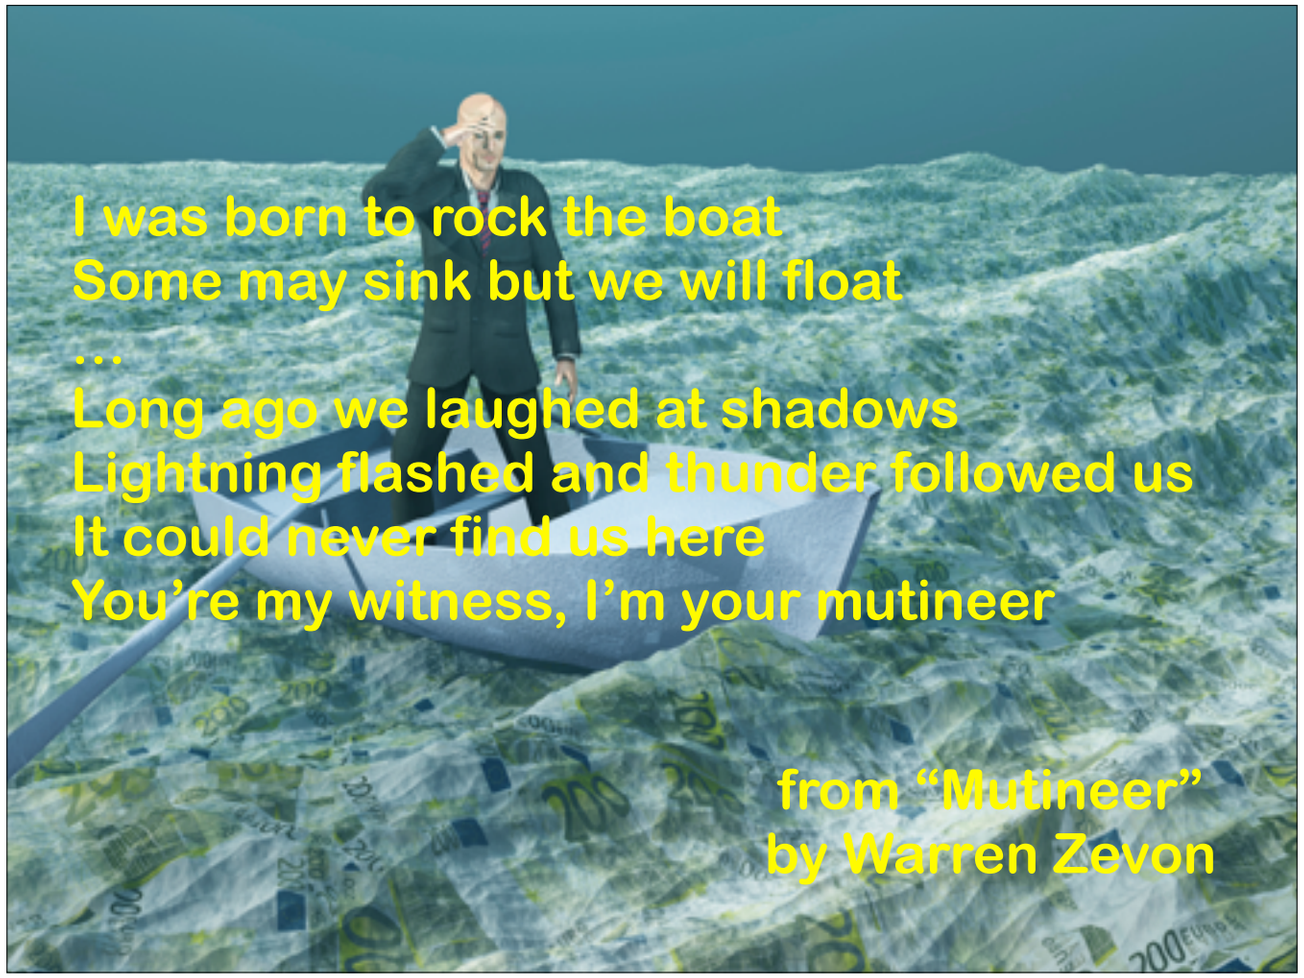
I was born to rock the boat
Some may sink but we will float
...
Long ago we laughed at shadows
Lightning flashed and thunder followed us
It could never find us here
You're my witness, I'm your mutineer
from "Mutineer"
by Warren Zevon

To Do

- Find and use a word for the rock-the-boatness of markets
- Learn how to tax market pollution
- Play

Questions?

Dale Rosenthal proposed the I tribe — as in Inattentive. If they were always inattentive, then they would tend to be a force for lowering volatility. However, on the way to Jak's afterwards, there was the observation that while they would usually be inattentive, they would pay attention to big news. Thus they might be a force for increasing volatility rather than decreasing it.

David Smith asked about news being distorted more on the downside that the upside. I stumbled through an answer. What I attempted to say was that we could find lots of ways to have a stronger distortion on the downside, but finding reasonable ones could be hard. At the time I was confounding distortion of news with asymmetry in utility — they are related but I think not the same.

There was a comment/question about agents changing their minds. It is my understanding (from what little I know about agent-based modeling) that mind-changing is an important feature of crashes.

Finally there was: "So what have we learned about taxing strategies?" My answer was that the difficulty of the task has been confirmed.

A question that I wish had been asked:

Why did you remain willfully ignorant of agent-based modeling?

There is the practical answer that Dirk would have chased me for miles if I had spent my time studying rather than coming up with a talk.

But there is a more philosophical reason as well. The “Play” item in the final version of the to-do list is a serious entry, and here is why:

The marketAgent package is independent of other agent-based models. Probably it is naive and makes mistakes that others have run across and eliminated. But there is a very small chance that it does something new and useful. Always studying what has gone before means our effective sample size on how to do the task is close to 1.

This is how our school systems are run. Children enter school as explorers, but that is struck out of them pretty sharply. Only in graduate school is exploring encouraged. I think that is completely backwards — exploring should be always on, and graduate school should be where complete knowledge is consolidated.

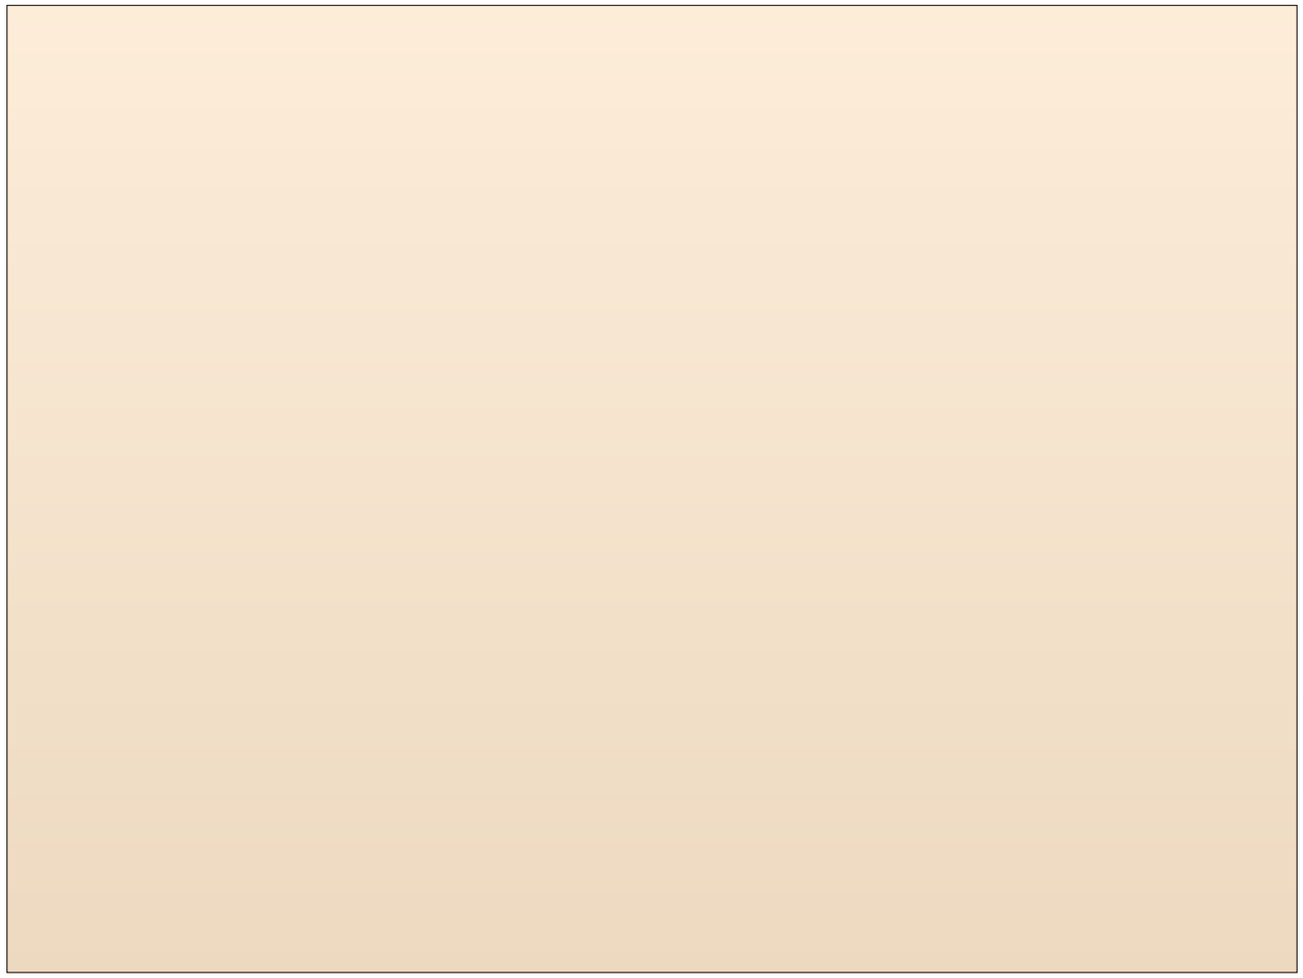

Another question that was not asked:

So are your calls to doRun verbs or nouns?

Definitely verbs, but the package could do with a dose or two of Rcpp.

The best question was asked at Jak's post-conference:

What fraction of days are you SBS versus SBU?

My answer: I'm SBS everyday, occasionally I'm SBU.

> … the sense of precarious dependence of all we know upon linguistic tools which themselves are largely unknown need not be discouraging to science but should, rather, foster that humility which accompanies the true scientific spirit, and thus forbid that arrogance of the mind which hinders real scientific curiosity and detachment.
>
> — “Science and Linguistics” by Benjamin Lee Whorf

For this version, I’ll leave the last word to Benjamin.