Information Cascades, starring Bikchandani, Hirshleifer, & Welch 1992

Kevin Dorst SLLERG

I. The Big Question

Human society characterized by *global diversity* but *local conformity*: what we wear; how we talk; where we eat; what we read; which bike lights we buy; which candidates we vote for; what we value; what we believe; how we think.

I'll be drawing centrally on Bikchandani, Hirshleifer, and Welch 1992.

Sometimes this is harmless or good. Sometimes it's quite bad.
→ Big social engineering problem!

Cf. the 2016 presidential election cycle.

Many explanations: social sanctions; coordination benefits; preferences for conformity ("peer pressure"); communication.

Big Question: To what extent might this global-diversity-plus-local-conformity be due to *epistemically rational* social influences?

Help answer *why/how* people are influenced by others, and spark ideas for how we can deal with it.

II. Information Cascades

BHW's idea: people are influenced by the information conveyed by others' actions. Sometimes this can lead widespread to conformity based on very little information. Such conformity is *fragile* and can be (in a sense) *accidental*.

So maybe explain (rational!) global diversity + local conformity.

Information Cascade: When it is optimal for an individual, having observed the actions of those ahead of her, to follow the behavior of the preceding individual without regard to her own private information. (BHW: 992)

Question: should "When *it is common knowledge that* begin this definition?"

Example: marbles!

Common knowledge that 50-50 chance between majority ($\frac{2}{3}$) blue and majority ($\frac{2}{3}$) green, and that people will guess the majority color they (rationally) think more likely.

It's **common knowledge** that p iff everyone knows p, everyone knows that everyone knows p, and so on. This will (maybe) be important.

1. Ada draws; guesses majority blue. Everyone can infer she *drew* blue, so that fact becomes common knowledge.

2. Bob draws; guesses majority blue. Everyone infers that *he* drew blue too.

Suppose common knowledge that if someone is 50-50, they'll guess the color they drew.

3. What about Cherie? *No matter the color of her draw, she'll guess majority blue.* She is in an information cascade. Everyone knows this, so her announcement is uninformative.

4, 5,... Therefore Dillon is in the same epistemic position as Cherie. So *his* guess of majority blue is uninformative too. Likewise for Eve, Fred, and all the rest. They are cascading.

Notice a few things:

Easley & Kleinberg 2010, Ch 16

1. Here imitation is epistemically rational—no peer pressure required.

2. Given the right setup, cascades can occur very easily: probability that Cherie will be in a *correct* cascade is $\frac{4}{9}$; *incorrect* cascade is $\frac{1}{9}$.

We'll come back to how big this 'given' clause is.

3. Information cascades can fairly easily lead to radically non-optimal outcomes.

   Suppose majority green. Then $\frac{1}{9}$ probability that *everyone* gets it wrong (no matter how many people).

   This despite the fact that if they *ignored* everyone else's guesses, then as the group grows the probability that $\frac{2}{3}$ of players would guess correctly goes to 1.(And probability that *everyone* gets it wrong would go to 0 quite quickly; with $n = 6$, it is $(\frac{1}{3^5}) \approx .004$.)

What's going on? *Epistemic* tragedy of the commons. (Baltag et al. 2013).

4. *Cascades are fragile*. For they are (commonly known to be) based on very little information—in our case, two marble draws. This means small shocks can shatter cascades.

   Public release of information.

   More informed individuals ("experts") who can draw twice.

   Changes in payoffs.

Suppose cascading till Yousef, who displays his marble to room. Then Zina might "flip" the cascade and guess majority green.

Contrasts with social sanctions, coordination benefits, and conformity preferences (purely epistemic, fragile).
Contrasts with reasoning/communication (big errors likely).

III. Generalizing

Basic features for information cascades:

Easley & Kleinberg 2010

1. A binary [?] decision: accept/reject; assert/be silent; vote for/against...
2. Decisions made in sequence.
3. Each person has some (inconclusive) private information.
4. Cannot observe others' private information, but can see what they *do*.

Or some summary statistic of what they did.

If this is right, information cascades have a *huge* domain of application, and so represent a big social risk:

- Politics. "Momentum" (polls; Super Tuesday); group homogeneity → group polarization. "Lock her up"; "Build a wall."

Sunstein 2009; Anderson

- Group deliberation. Informal opinion-voicing on committees (Sunstein 2006);
- Hiring. "Rising stars."
- Economics. Financial booms, busts, and bubbles; bike lights.
- Sociology. Fads; fashions; culture??

BHW's 1992 model assumes the following: it is common knowledge that...

1. Binary decision; fixed linear order, $1, 2, ..., N$
2. All agents have same cost C and possible gain V of adopting.
3. Everyone has a common prior probability μ.

C fixed; V has possible values $v_1, ..., v_s$ with $v_1 < C < v_s$.

4. Each player n receives an inconclusive private signal X_n, with possible values $x_1, ..., x_r$. $(\forall X_n, x_i, v_j : \mu(X_n = x_i | V = v_j) > 0)$

X_i are independent and identically distributed (wrt μ, conditional on V).

5. Each agent a Bayesian reasoner who maximizes expected utility.

6. Higher x_i values make higher v_i values more likely.

Agent n is in an **information cascade** iff she accepts (rejects) no matter what her private signal is.

Theorem (BHW 1992). As number of individuals increases, the probability that a cascade eventually starts goes to one.

Easley & Kleinberg 2010: There are many ways to loosen the model of reasoning, but they lead to "qualitatively similar conclusions."

Baltag et al. 2013: If you *tighten* the model by formalizing agents' higher-order reasoning, information cascades are still possible. Cascades can't be solved with individual rationality.

Even recognizing *that* they are in a cascade won't help.

Following Anderson (2017 MIT colloquim), one salient feature of politics is that people are not just truth-motivated. Often assertions ("Lock her up"; "Build a wall"; "Trump's inauguration crowd was bigger than Obama's") are *identity-expressive*, voicing solidarity with a group. Proposal: this can exacerbate the problem. What if we *don't know* whether people are trying to get to the truth, or merely following a crowd?

Preliminary results: looks like information cascades become *more* likely.

IV. Concerns

General concerns: models include massive common knowledge assumptions. And (1) such common knowledge doesn't seem to hold in many applications, (2) there's reason to worry that we *never* have common knowledge (Lederman 2017b), and (3) often dropping such common knowledge assumptions breaks simple game-theoretic reasoning (Lederman 2015, 2017a).

So what happens when we drop various common knowledge assumptions?

Specific concern: Do genuine information cascades *ever* happen in real life? Note that definition requires that you accept (reject) *no matter your private information*, i.e. given any information-state you could start with.

Q: How broad is scope of this 'could'?

In their definitions, $\{x_1, ..., x_r\}$ correspond to several things you could mean by "Ada's possible information":

- The signals that Ada thinks she could have received.
- The signals that Bob thinks he could have received.
- The signals that Bob thinks Ada might have received.
- The signals that Cherie thinks Bob might think that Ada might have received
- The signals that Bob thinks *someone* might have received.
- The signals that *anyone* thinks *someone* might have received.

I think this last one is what is needed for *fragile conformity*. Let X be the set of possibilities Cherie "could" have received.

> Suppose X is (no larger than) the set of signals that anyone leaves open *Cherie* might have received. Then her action is uninformative to Dillon. But he might have a bigger set of signals, so *he* might not follow her.

No guarantee of conformity.

> Suppose X is the set of signals that anyone leaves open *someone* might have received. Suppose by Cherie's turn, any signal in X would lead her to A. Then any signal *Dillon* leaves open she might have would lead her to A, so her action is uninformative to him. Since he also has a signal in X, he also does A. And so on.

Surprise (for me)! I don't think we need X to be broader, including all the signals they *commonly* leave open—i.e. those such that someone thinks maybe someone received OR someone thinks maybe someone thinks maybe someone received, OR... [and so on]. Suppose X is as before, and no bigger. Then Eve might not know that Dillon knew that Cherie would do A. But she did know that (i) if Dillon went during Cherie's turn, Dillon would do A (any signal in X would lead to A), and (ii) Cherie's doing A won't provide Dillon with reason *not* to do A; so she can infer (iii) that Dillon will do A.
(Q: Is this right? Does the reasoning that led up to Cherie doing A matter? Does it matter that they won't necessarily know what X is?)

This X is pretty broad. If *anyone* leave open that *someone* got a signal that wouldn't be overridden by the public choices, then—even if it was common knowledge that *Cherie* wouldn't get such a signal—she's not yet cascading.

Bike lights aren't like that. Nor is voting, following your group, voicing an opinion on a committee, believing Clinton ran a sex-slave ring, etc.

Yet we still want to say there is an important sense in which these people are ignoring their own information and following the crowd. Potential weaker notion of interest:

> Agent n is on an **information slide** if given just her private information she'd accept (reject), but adding the public information leads her to reject (accept).

Modify this to include mutual or common knowledge/likelihood?

Information slides don't have the simple (fragile conformity) properties of information cascades. But they seem much more realistic/common.

Could they have similar effects? Could the possibility of information slides also increase the chance of radical group error as compared to ignoring public information?