STIeT Socio-Technical Infrastructure for Electronic Transactions

Spam Economics, Bonds, and Restoring Valuable Communication

Thede Loder, Marshall Van Alstyne,
Rick Wash, Mark Benerofe

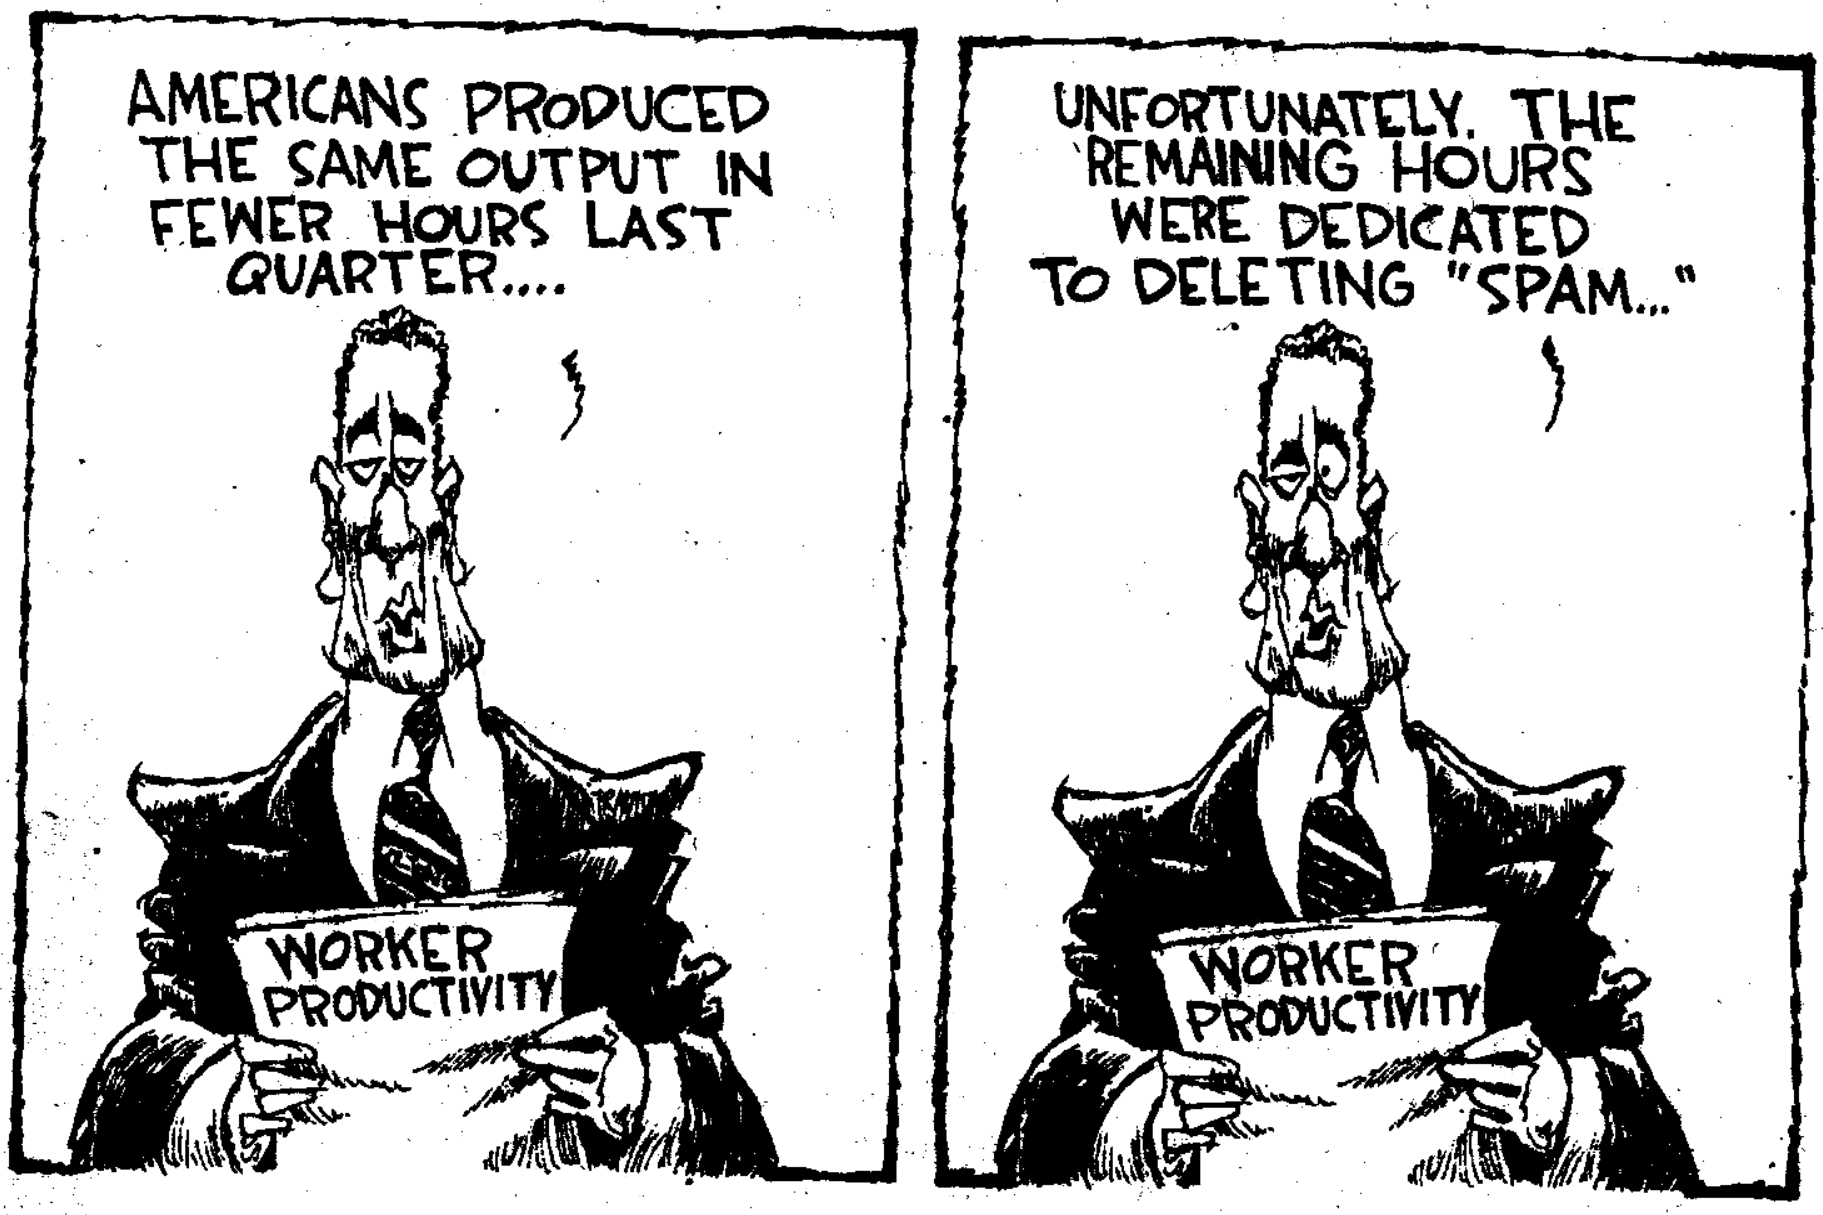

Nick Anderson
The Louisville Courier-Journal
Washington Post Writers Group

Data on the Spam Problem

- Estimated yearly loss to US businesses $10 billion (Ferris Rsch)
- ISPs estimate the cost of spam at $2-$3 per user per month (IDG)
- The Federal 'CAN-SPAM' Act passed with near unanimous support; superceding 8 State laws, 6+ spam laws still pending in congress (NYT)
- ≥ 60% of all email is now spam (Brightmail)
- 29% of Americans report curtailing e-mail use (Pew Internet report)

Spammer Alan Ralsky vows to carry on (NYT)

Existing or Proposed Solutions

- Legislative/Regulatory:
  - Banning, labeling, opt-out
  - Taxation
- Technological:
  - Filtering: Rule based (static or dynamic), Naïve Bayesian Filters, collective/community classification
  - Challenge-Response: CAPTCHAs (reverse-Turing tests, return address testing
- Economic
  - Computational challenge (proof of work)
  - E-stamps, sender-pays

Problems with Legislative Solutions

- Defining spam
  - Whose definition?
  - One size fits all, disregards individual preferences
- Regulating: banning and labeling
  - Enforceability, jurisdiction?
  - Costly to police and adjudicate
  - Labeling lacks incentive compatibility
- Taxing: e-stamps
  - Blocks wanted along with unwanted email
  - Blunts cost-effectiveness of email as medium

Problems with Technology Solutions

- Filtering (Rules, Collaborative, Bayesian):
  - False positives, false negatives
  - Increases spam email traffic
  - Costly arms-race
  - Consensus definition
  - Blocks automated email
- Challenge-Response (Reverse Turing Tests)
  - Cheaply hire real people to take test
  - Deadweight loss of human time

Modeling the Message Value

- Each message has a party-dependent value:
  - value to the sender *s*
  - value to the recipient *r*
- Range of *s* and *r* bounded by $\underline{r}$, $\overline{r}$, $\underline{s}$, $\overline{s}$

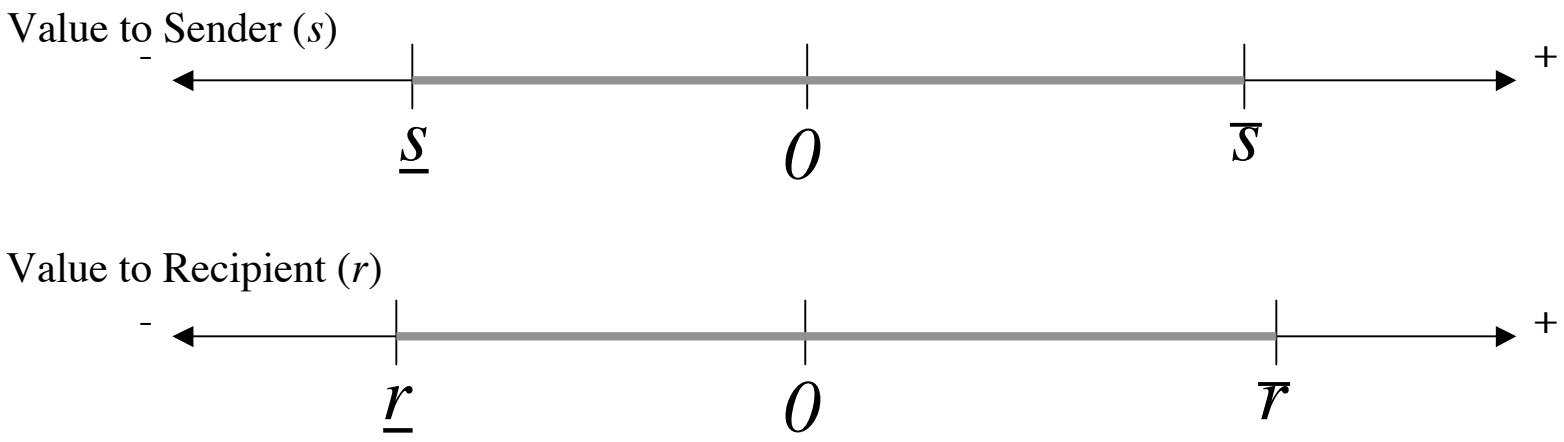

Model Assumptions

- Sender knows own value s and cost c_s *before* sending, won't send when $s < c_s$
- Sender does not know r (relaxed later with improved results)
- Recipient only knows her r value *after* reading and incurring cost c_r

Email Value Distribution

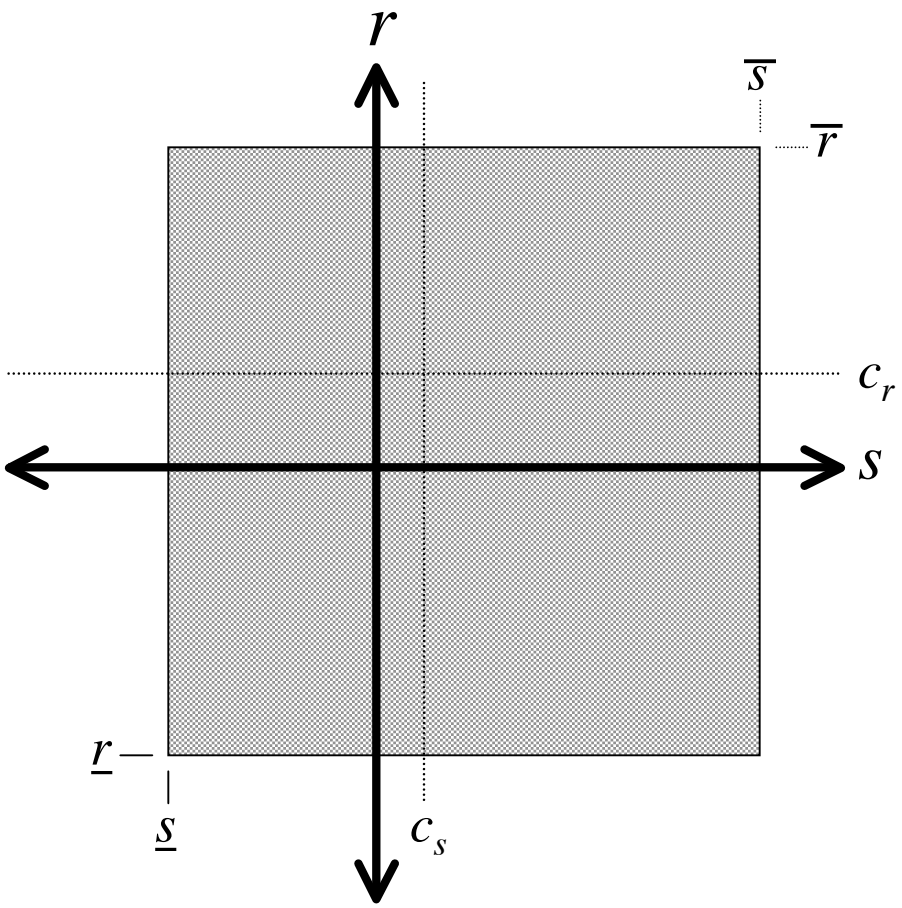

c_s = cost to send
c_r = cost to receive

region of positive probability

Email Value Distribution

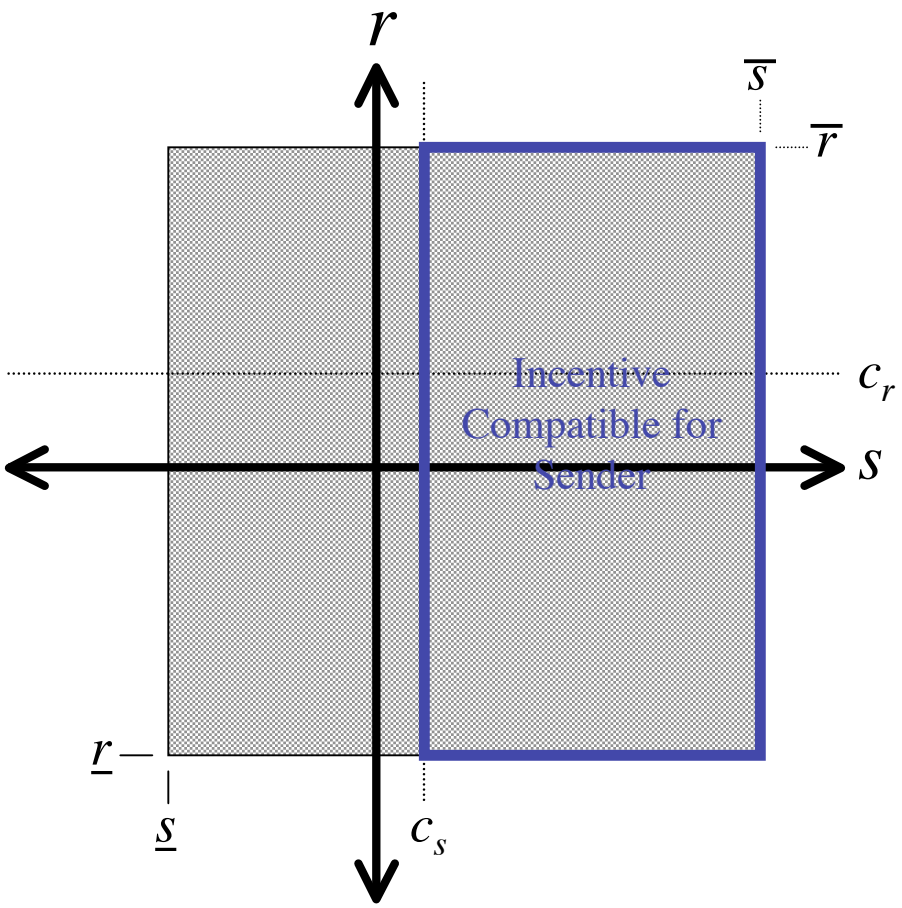

c_s = cost to send
c_r = cost to receive

- message is sent when $s > c_s$
- sender expects positive payoff

region of positive probability

Email Value Distribution

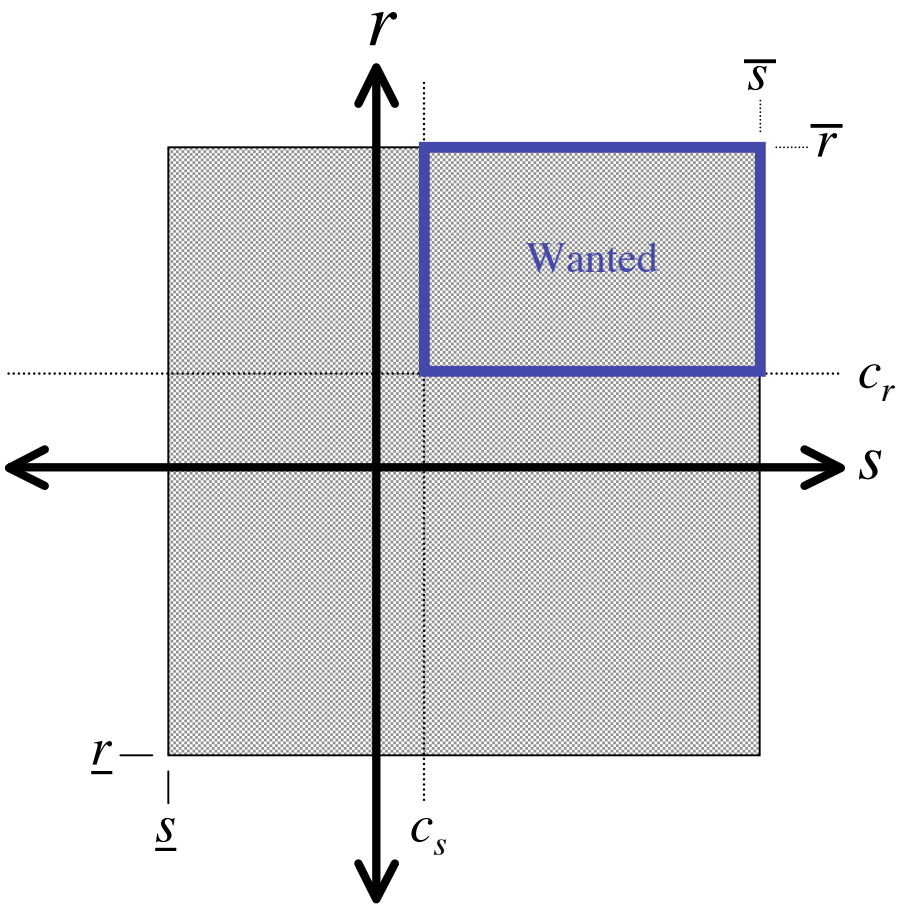

c_s = cost to send
c_r = cost to receive

Examples:
- project update from business associate
- long lost high school buddy says hi
- friend's wedding announcement
- shipping update from ecommerce vendor
- recruiting lead from a friend

region of positive probability

Email Value Distribution

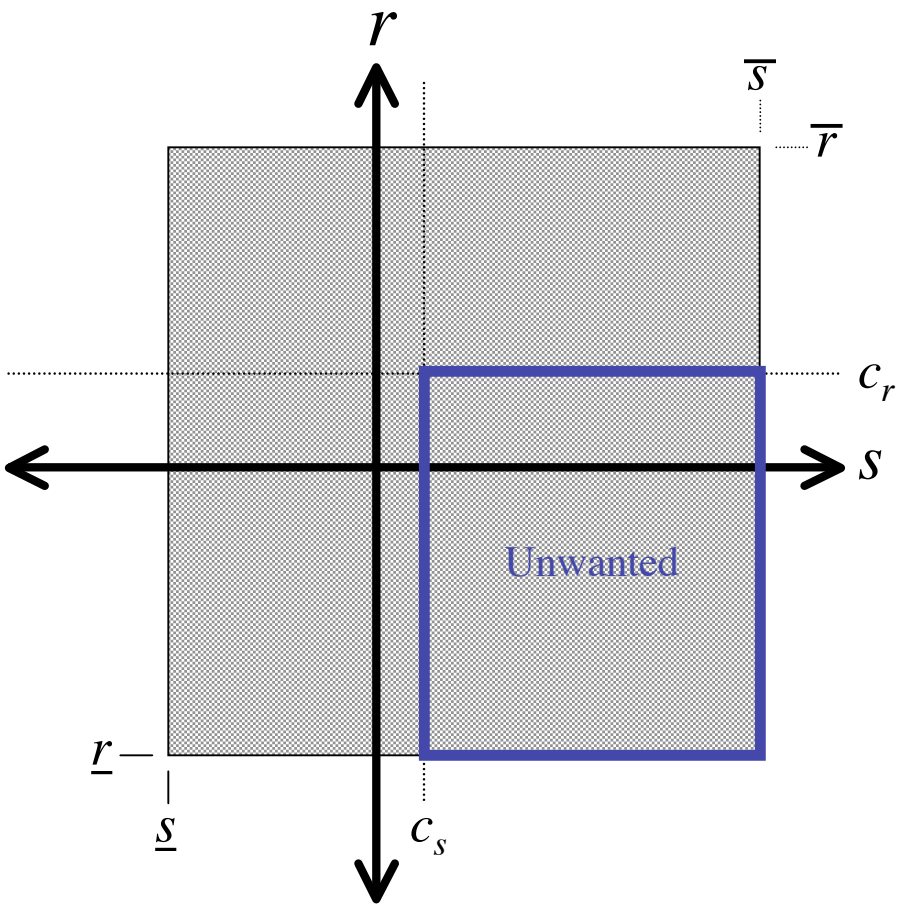

c_s = cost to send
c_r = cost to receive

Examples:
- offensive pornography
- message with attached virus
- phishing, Nigerian scam

region of positive probability

Email Value Distribution

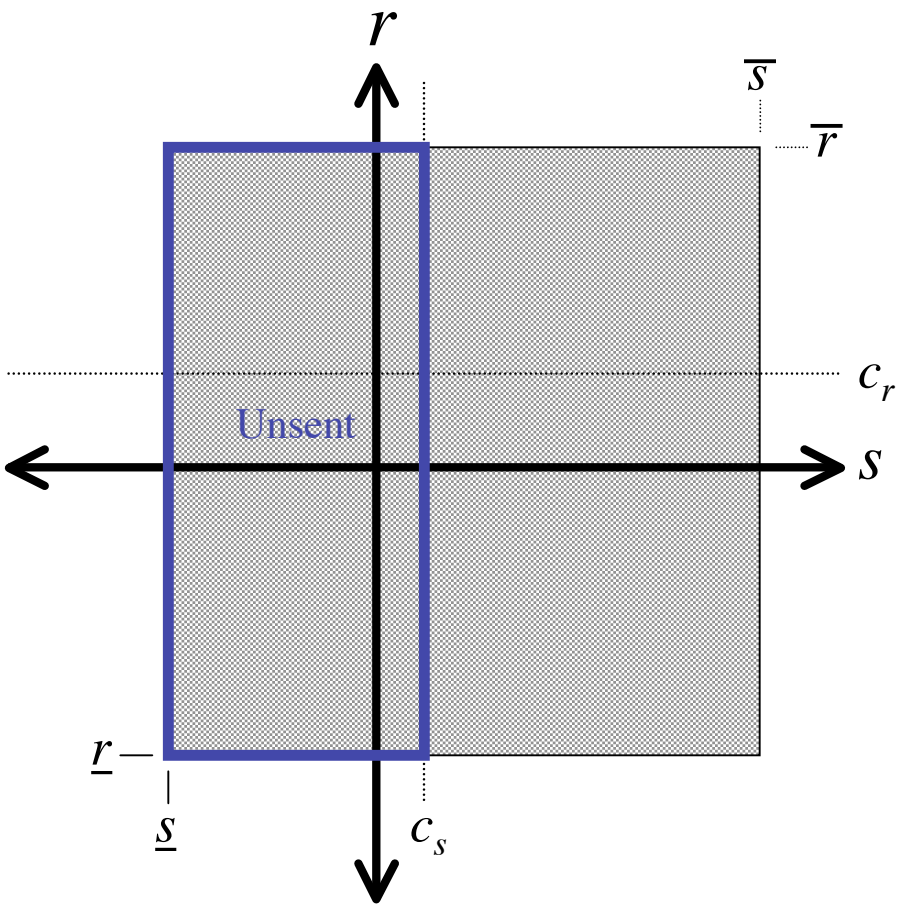

c_s = cost to send
c_r = cost to receive

- email unsent when $s < c_s$
- sender expects negative payoff

region of positive probability

Email Value Distribution

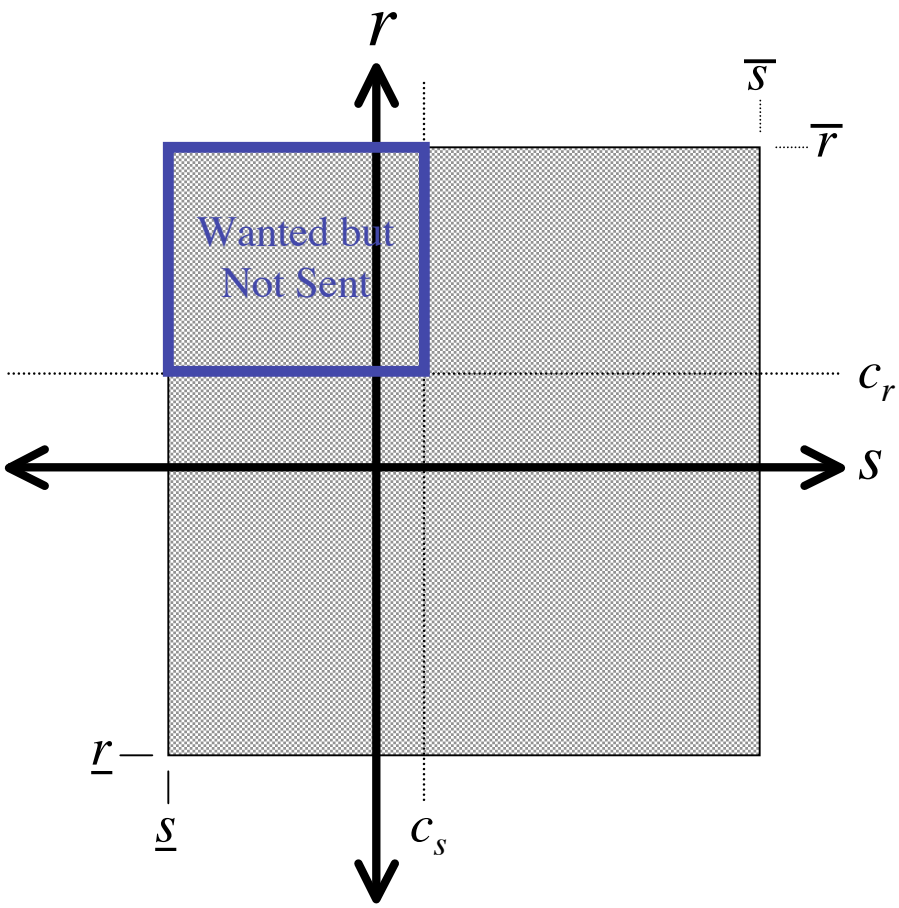

c_s = cost to send
c_r = cost to receive

Examples:
- personalized loan application
- customized news or subscription content
- costly sales leads

region of positive probability

Email Value Distribution

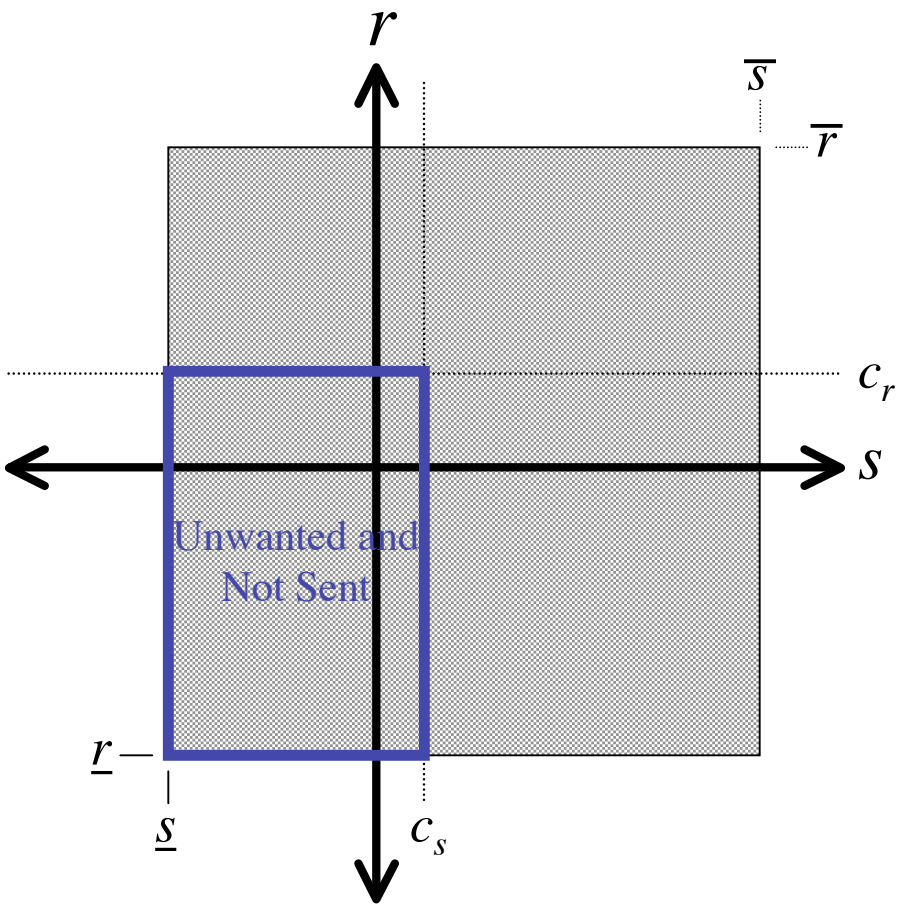

c_s = cost to send
c_r = cost to receive

Examples:
- embarrassing news sent to wrong person
- outdated or incorrect facts
- meaningless details

region of positive probability

Derivation of Payoffs

- Baseline Recipient Surplus :

$$RS_0 = \iint_A (r - c_r) f(s,r) ds dr$$

- Baseline Sender Surplus:

$$SS_0 = \iint_A (s - c_s) f(s,r) ds dr$$

- Total Welfare:

$$W_0 = RS_0 + SS_0$$

Social Welfare Contribution

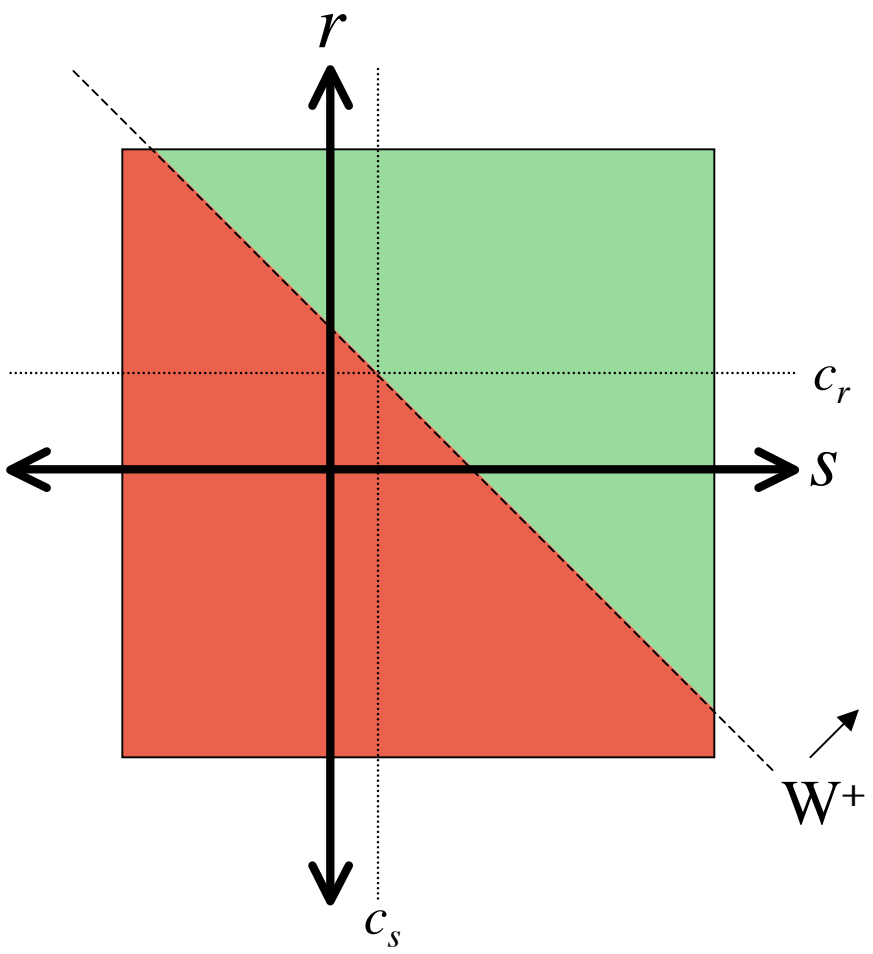

Welfare (W) = RS + SS

Received email northeast of the diagonal line makes a positive contribution to social welfare

positive contribution to W

negative contribution to W

Interpreting Existing Solutions

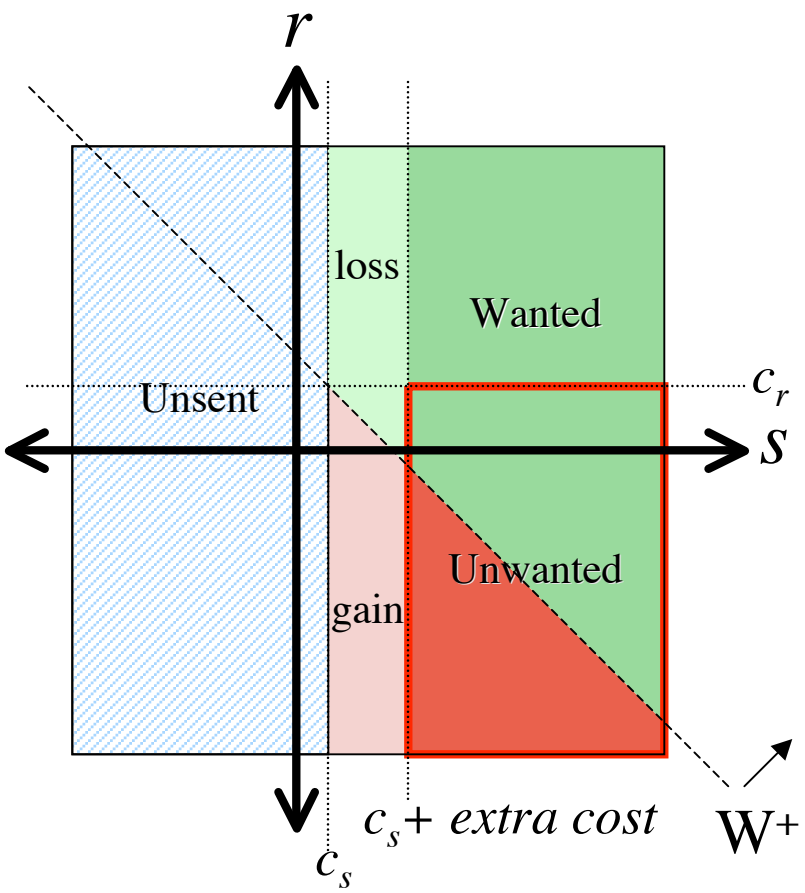

Solution	Sender cost	Social cost
Flat Tax	Per email tax	Tax collection, administration
E-stamps	Per email stamp	Stamp creation and collection
Computational Challenge	Compute cycles	deadweight loss, indiscriminate
Challenge Response or CAPTCHA	Human compute time	deadweight loss, indiscriminate
Criminalizing, Do-Not-Spam list	Evasion cost, fines	Enforcement Cost

Interpreting Filtering

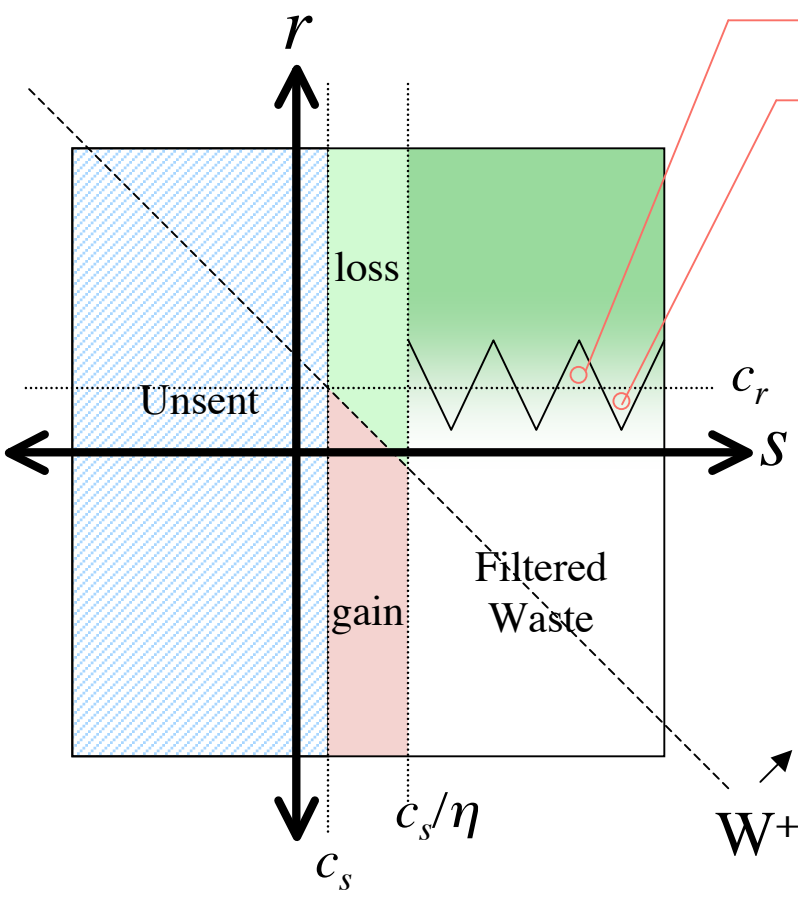

$\eta : 0 < \eta < 1$ email gets through
Sender decision:
send when $E(\text{gross return}) - \text{cost} > 0$

No filter
$E(\text{gross return})$ is s, cost is $c_s \rightarrow$
send when $s - c_s > 0$ or $s > c_s$

With filter
$E(\text{gross return})$ is ηs, cost is $c_s \rightarrow$
send when $\eta s - c_s > 0$ or $s > c_s/\eta$

An Economic Solution

Avoid Indiscriminate Defense

- Spam is different for each person:
  - "I can not define obscenity but I know it when I see it" -- Justice Potter Stewart

- Individualized Definition: "any email received that, if given the choice and advanced knowledge of its content, the recipient would have chosen not to receive"

"In terms of individual and aggregate social welfare, a system that facilitates valuable exchange and side payments will generally dominate a system that grants only unilateral veto power to either party."

Alternative: Attention Bonds

Initially, the sender knows more about the message than the receiver, so force them to reveal that private knowledge:

- Simple screening mechanism applied to *unrecognized* senders; known senders can be whitelisted with sufficiently strong authentication
- Challenge demands an escrowed bond fee of amount *b*
- Recipient has sole discretion to claim or return *b*.
- All proceeds *b* go to recipient.
- *Effects*:
  1. A recipient-controlled variable 'tax' on senders, based on sender behavior and message content
  2. Shift task from ex ante classification (hard) to ex post verification (easy).
- Implemented as the Attention Bond Mechanism (ABM)

Recipient Payoff with Bond

For any particular distribution, the recipient can remove the sender incentive to send email.

The sender sends when $s > cost$. Baseline cost is c_s. Any $s \mid s < c_s$ yields a negative expected payoff.

Positive payoff to recipient

Negative payoff to recipient

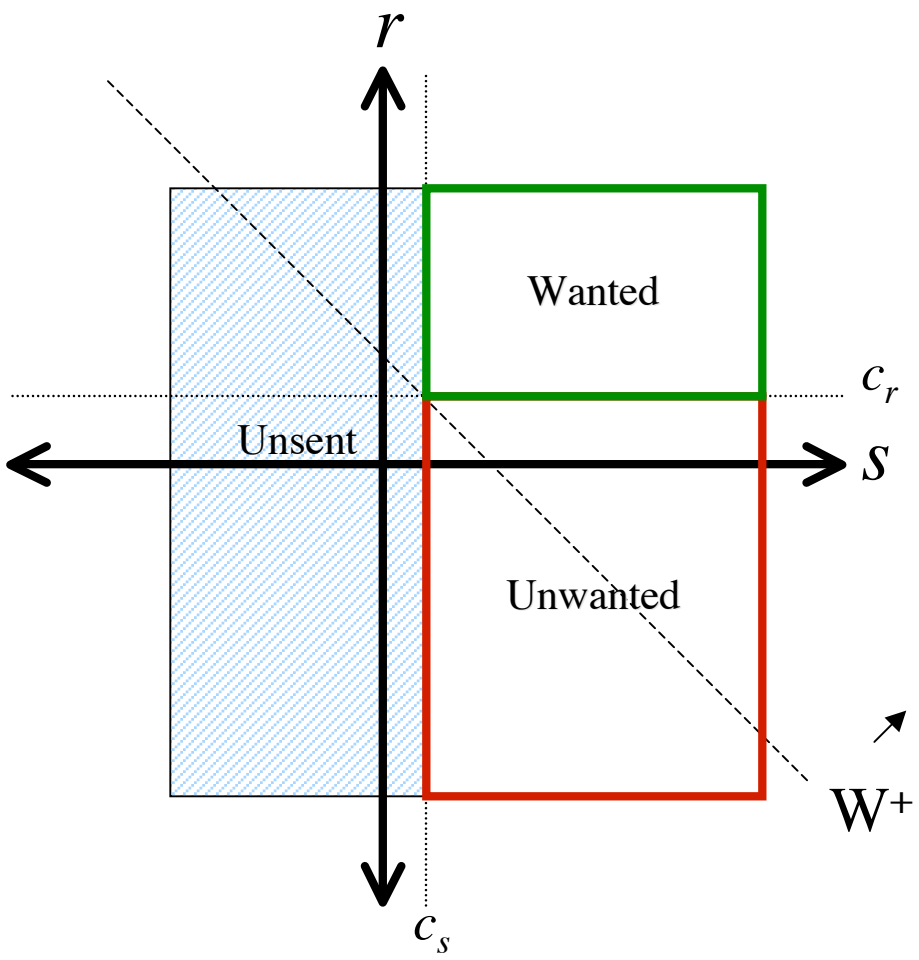

Recipient Payoff with Bond

When $b > 0$, sender costs are c_s+b

Increasing b eliminates incentive to send, blocking more of the email in the distribution

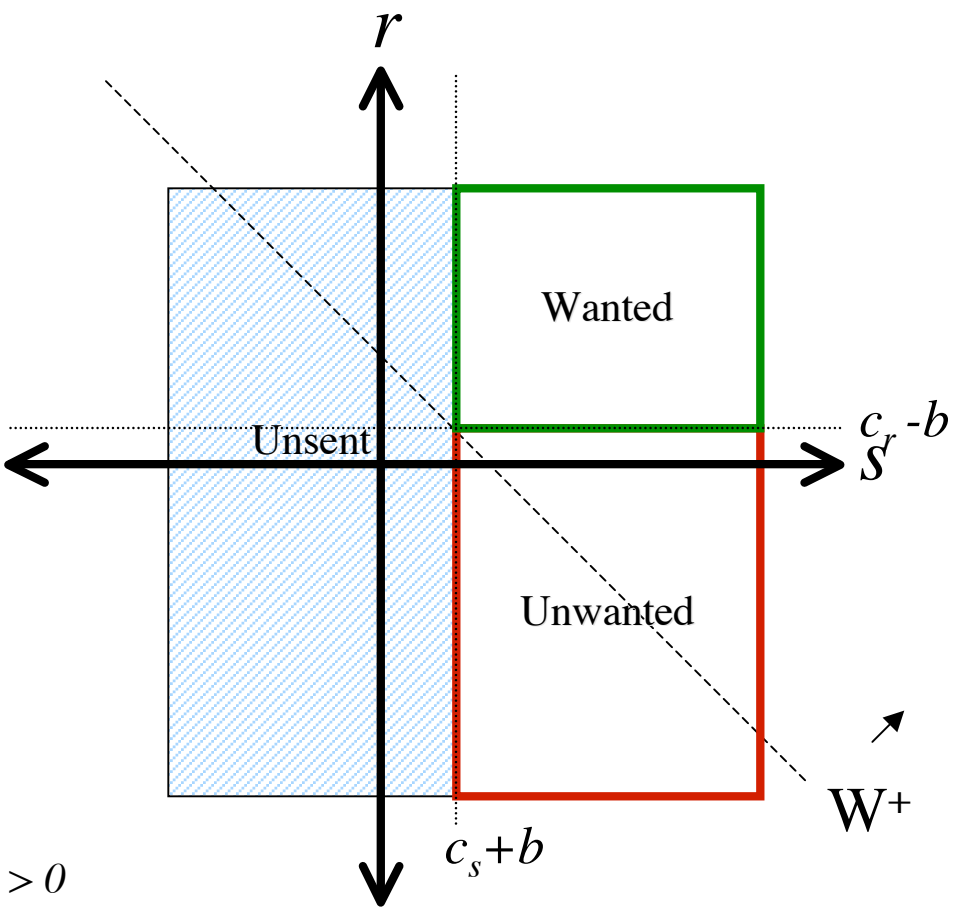

Positive payoff to recipient

Negative payoff to recipient

Now transfer b to recipient where $b > 0$

Recipient Payoff with Bond

When $b > 0$, sender costs are c_s+b

Increasing b eliminates incentive to send, blocking more of the email in the distribution

b increasing…

Positive payoff to recipient

Negative payoff to recipient

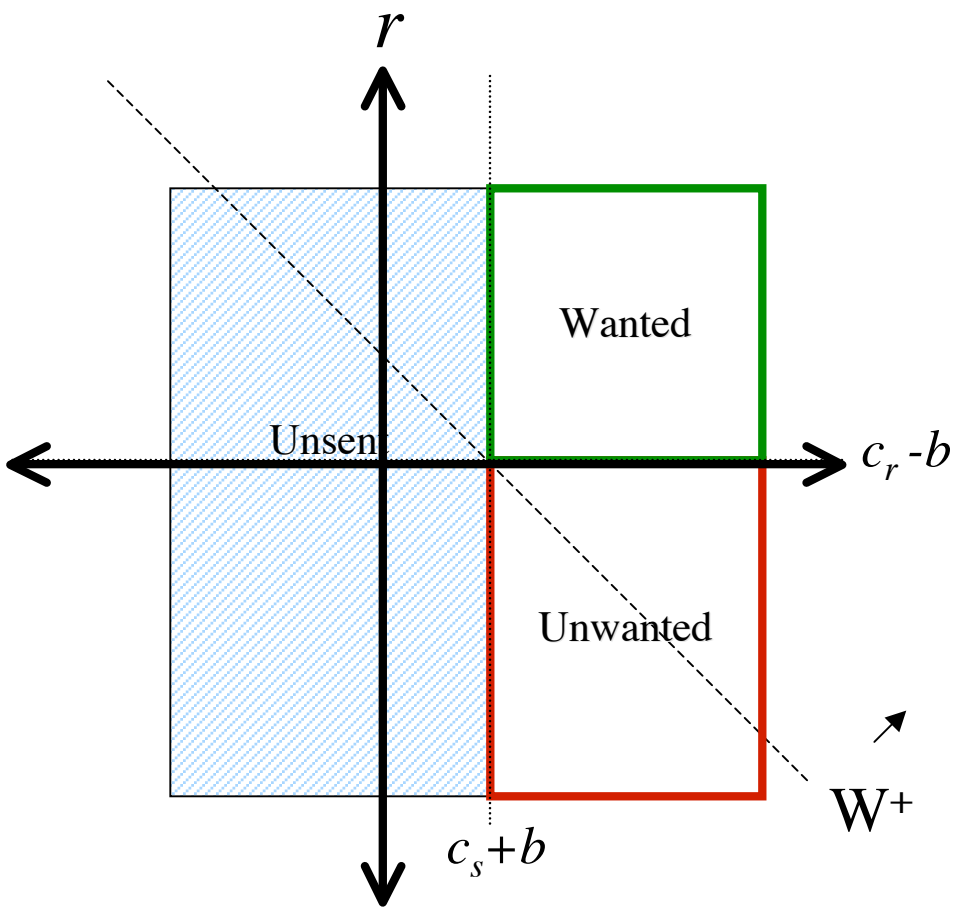

Recipient Payoff with Bond

When $b > 0$, sender costs are c_s+b

Increasing b eliminates incentive to send, blocking more of the email in the distribution

b increasing…

Positive payoff to recipient

Negative payoff to recipient

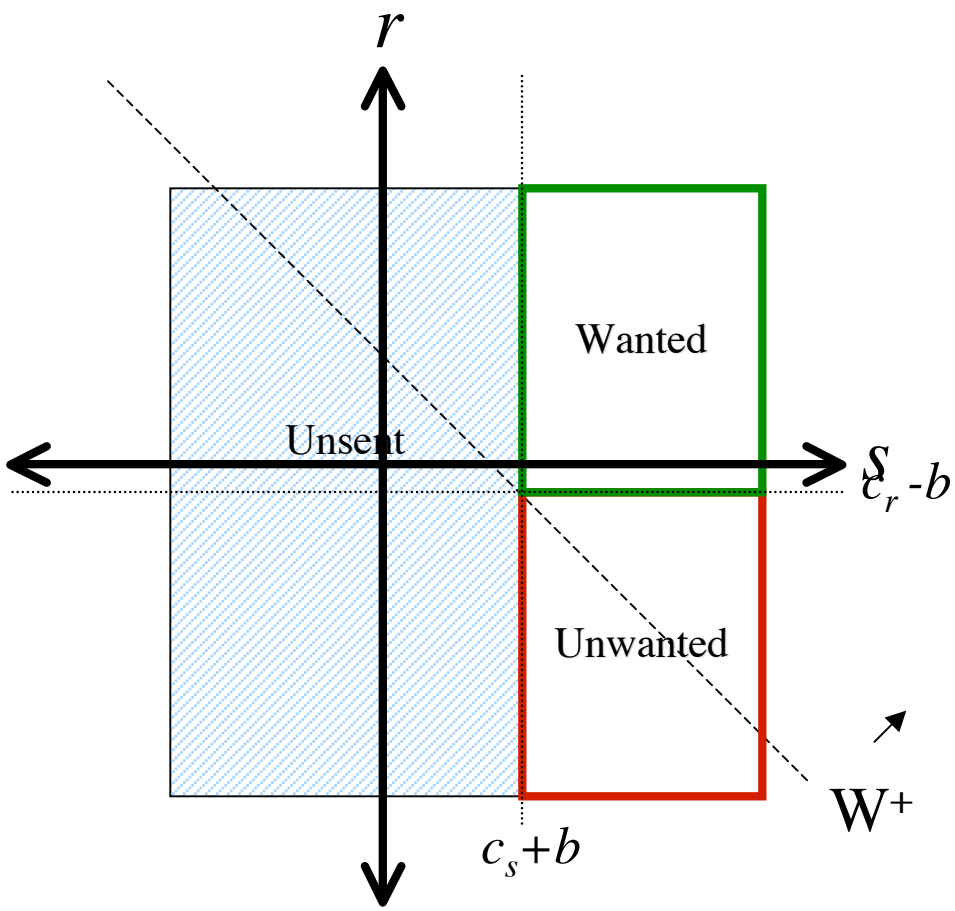

Recipient Payoff with Bond

When $b > 0$, sender costs are c_s+b

Increasing b eliminates incentive to send, blocking more of the email in the distribution

b increasing…

Positive payoff to recipient

Negative payoff to recipient

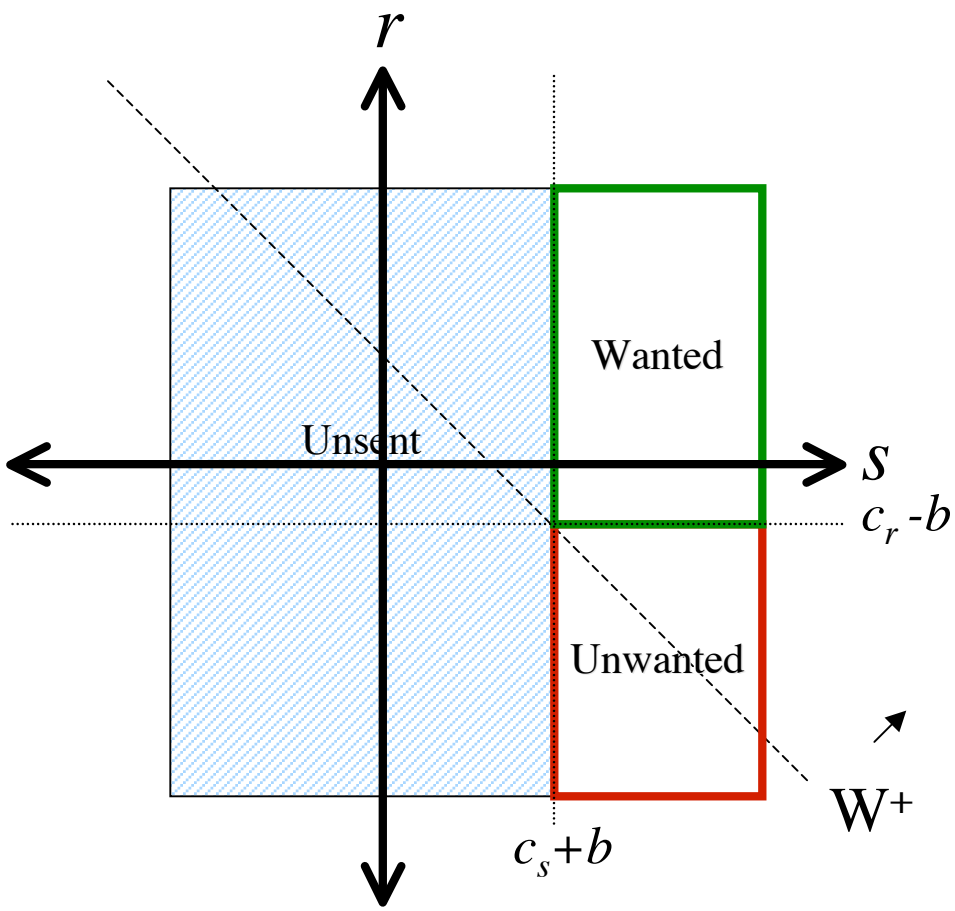

Recipient Payoff with Bond

When $b > 0$, sender costs are c_s+b

Increasing b eliminates incentive to send, blocking more of the email in the distribution

Positive payoff to recipient

Negative payoff to recipient

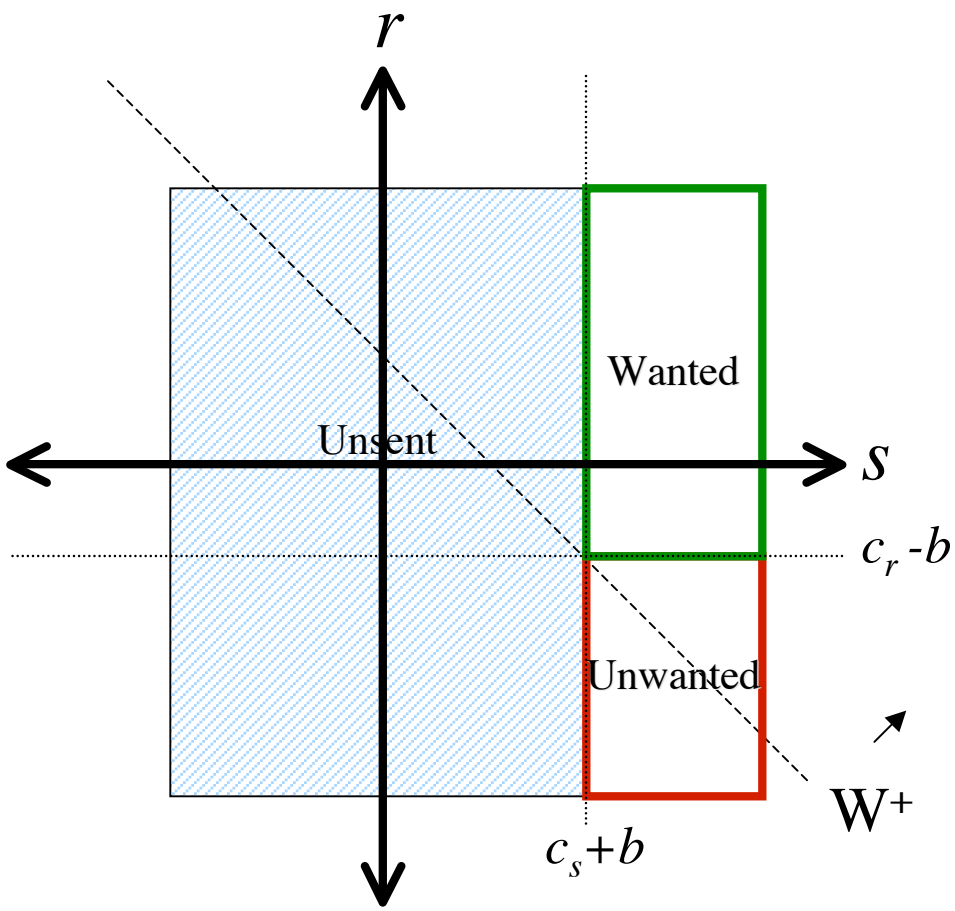

Recipient Payoff with Bond

When $b > 0$, sender costs are c_s+b

Increasing b eliminates incentive to send, blocking more of the email in the distribution

Positive payoff to recipient

Negative payoff to recipient

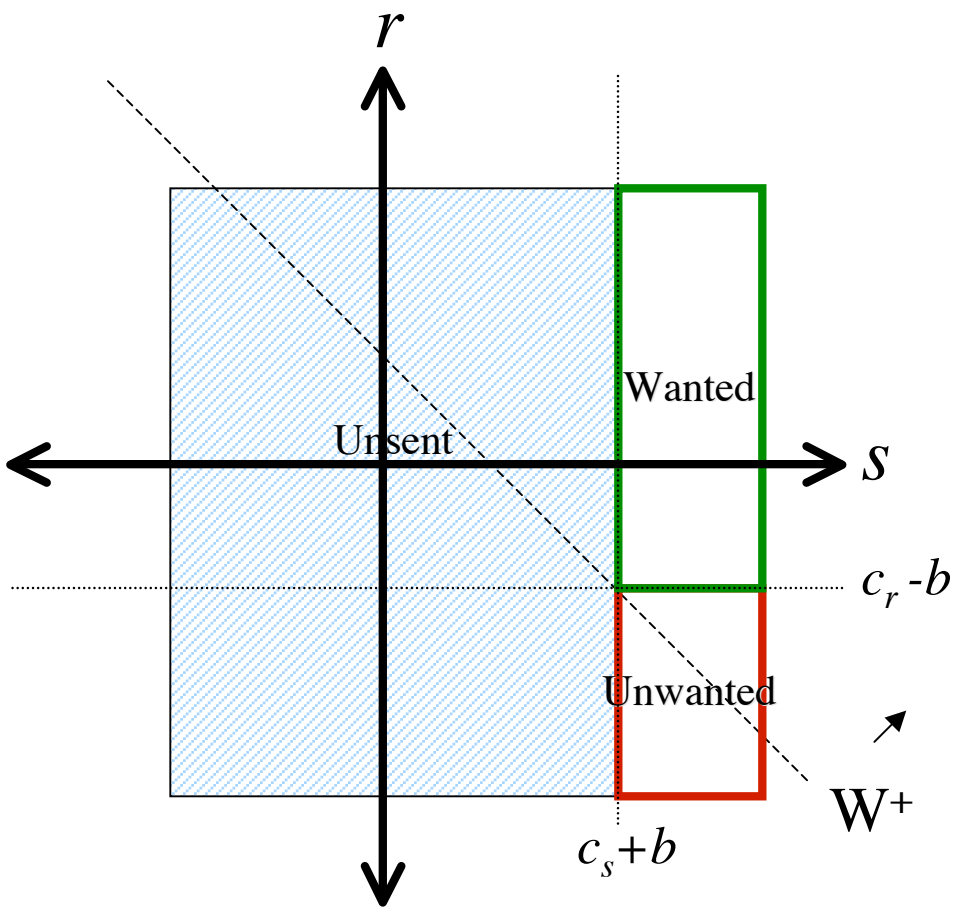

Recipient Payoff with Bond

When $b > 0$, sender costs are c_s+b

Increasing b eliminates incentive to send, blocking more of the email in the distribution

Positive payoff to recipient

Negative payoff to recipient

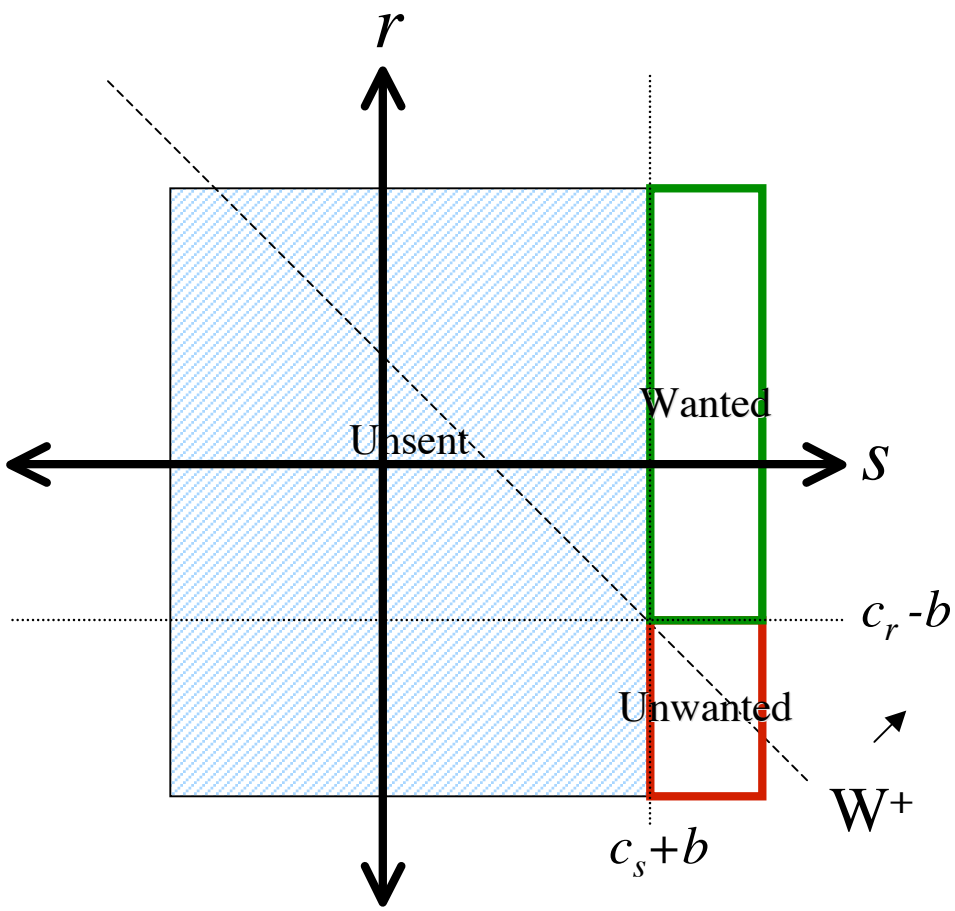

Recipient Payoff with Bond

When $b > 0$, sender costs are c_s+b

Increasing b eliminates incentive to send, blocking more of the email in the distribution

Positive payoff to recipient

Negative payoff to recipient

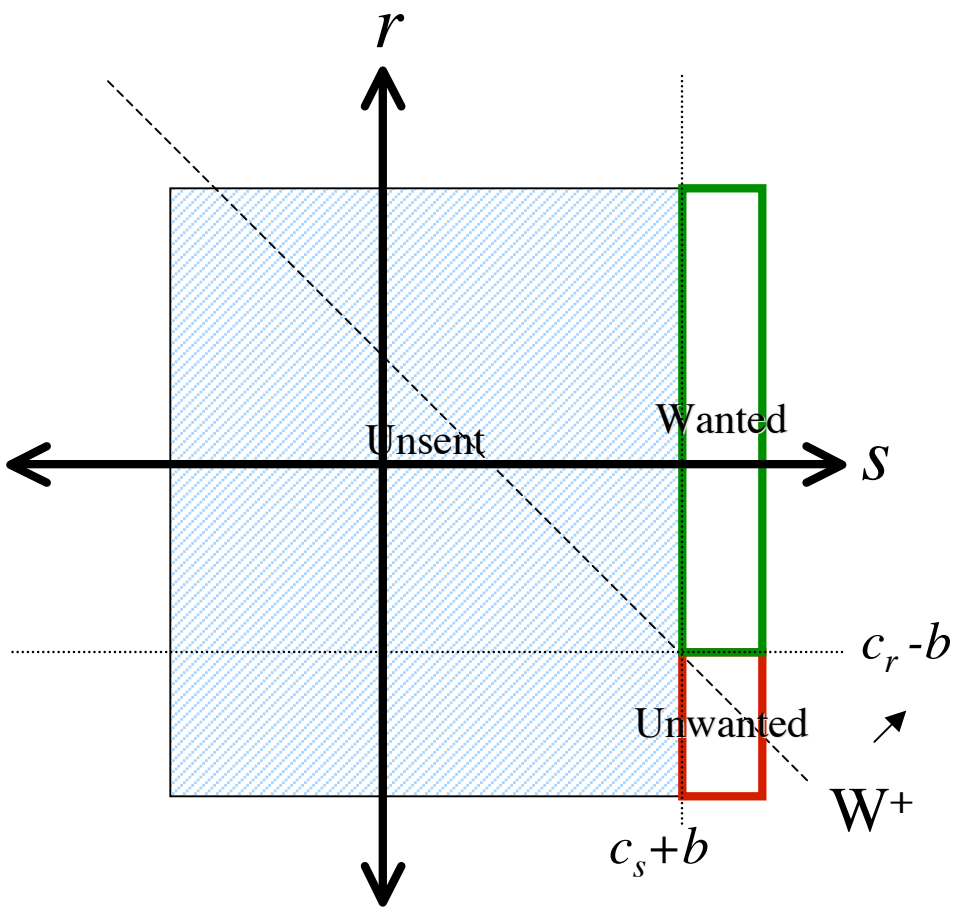

ABM Overview

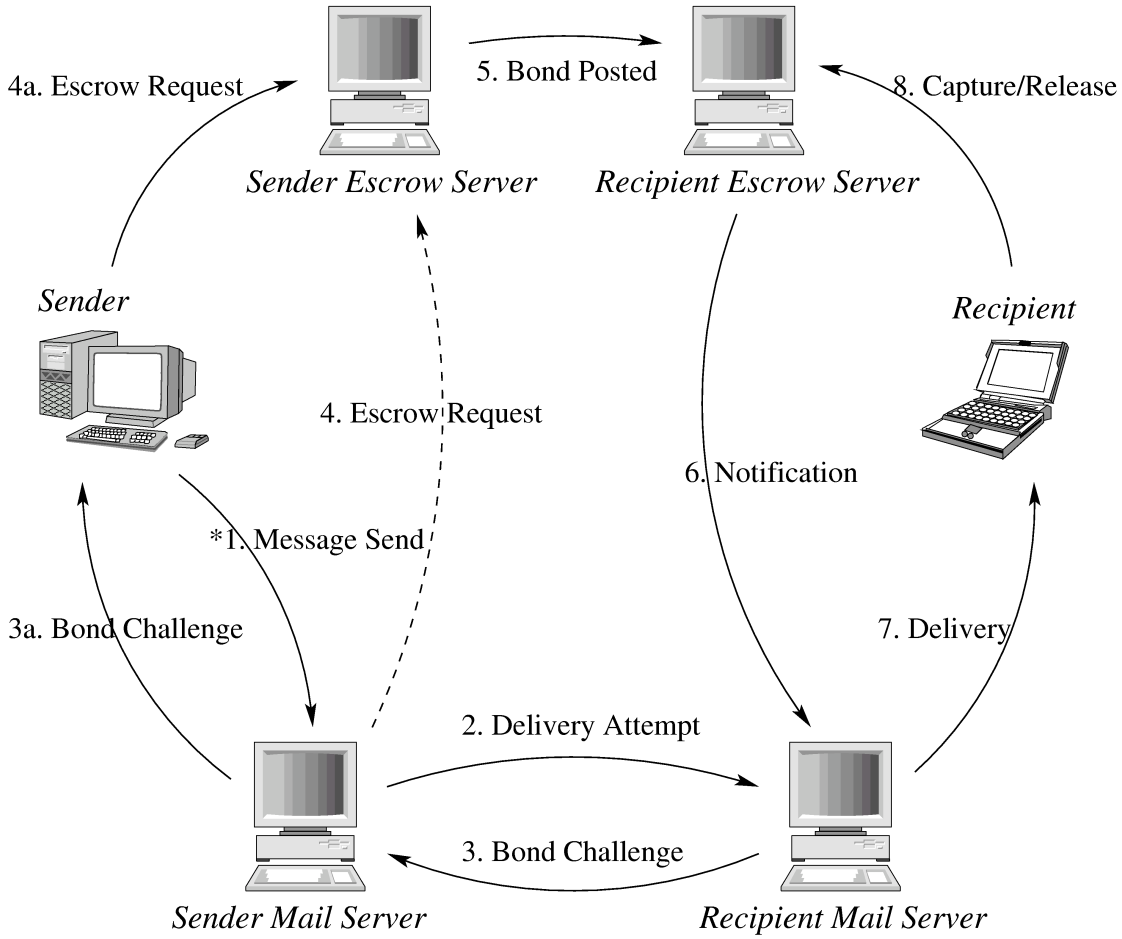

The ABM in Action

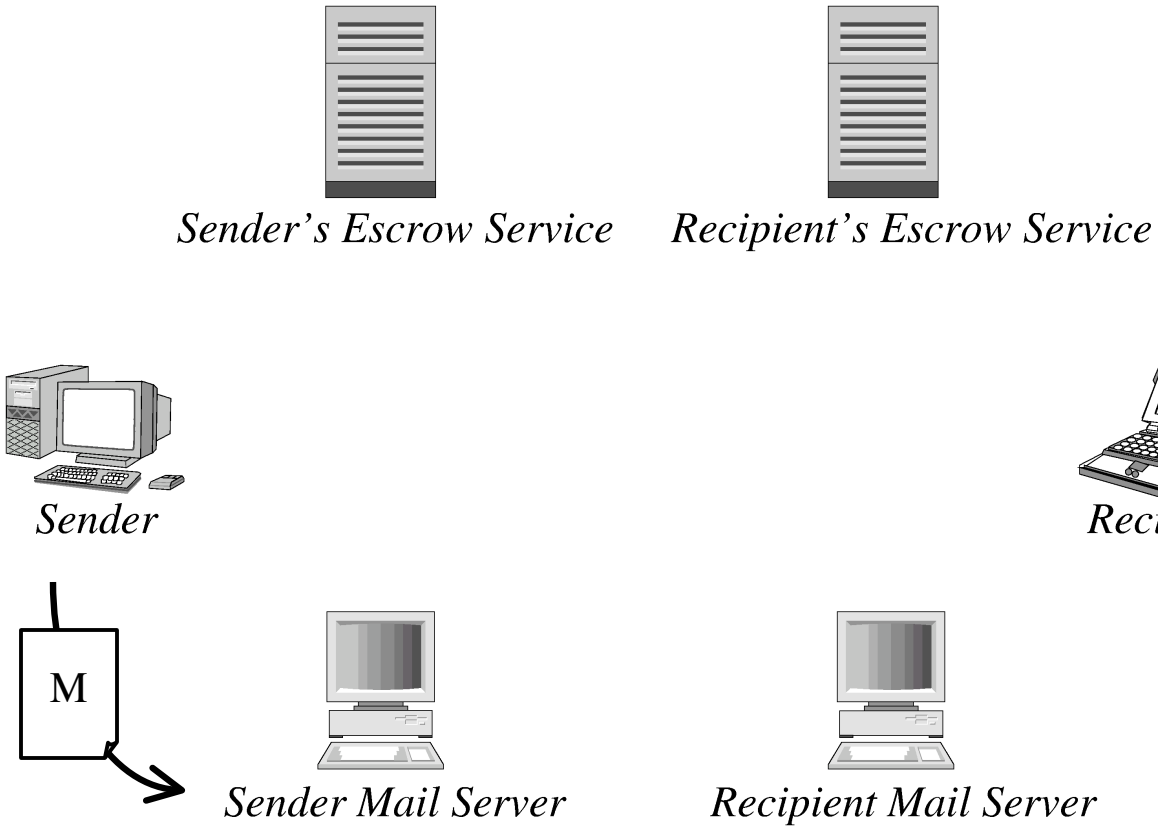

Sender sends a message

The ABM in Action

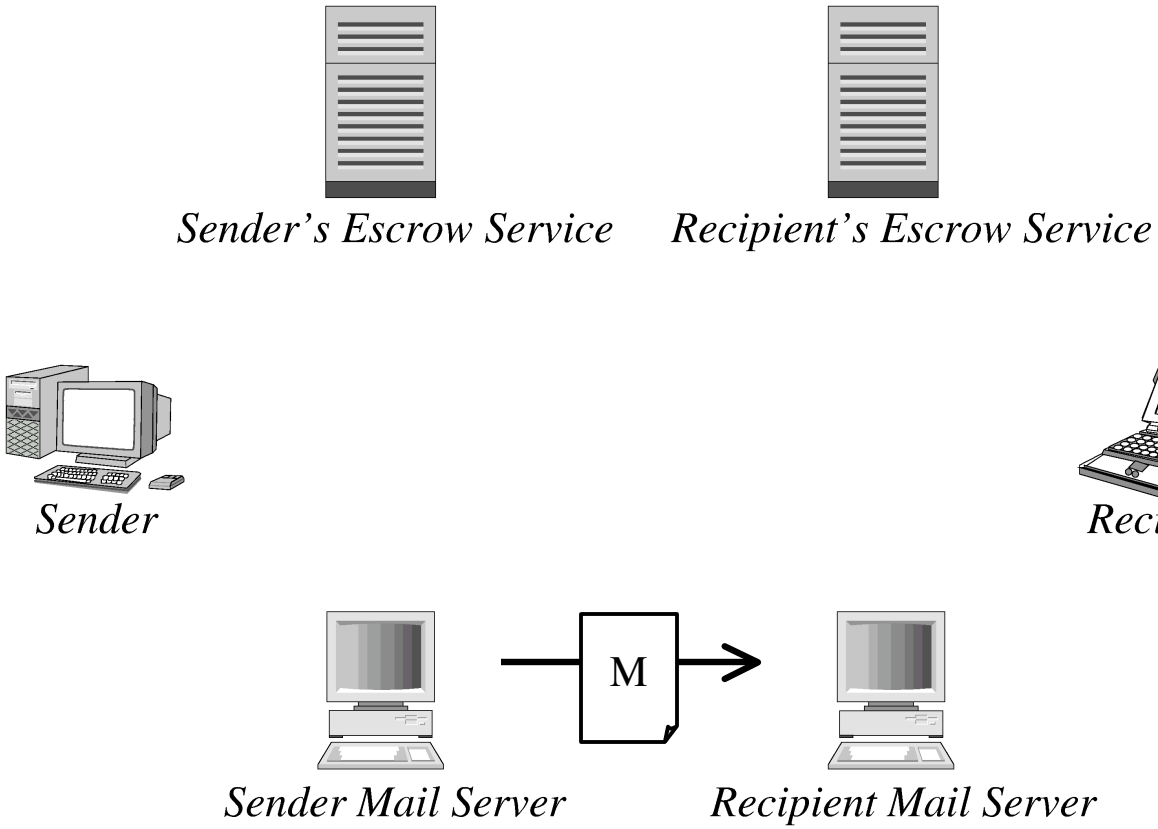

Sender mail server attempts delivery

The ABM in Action

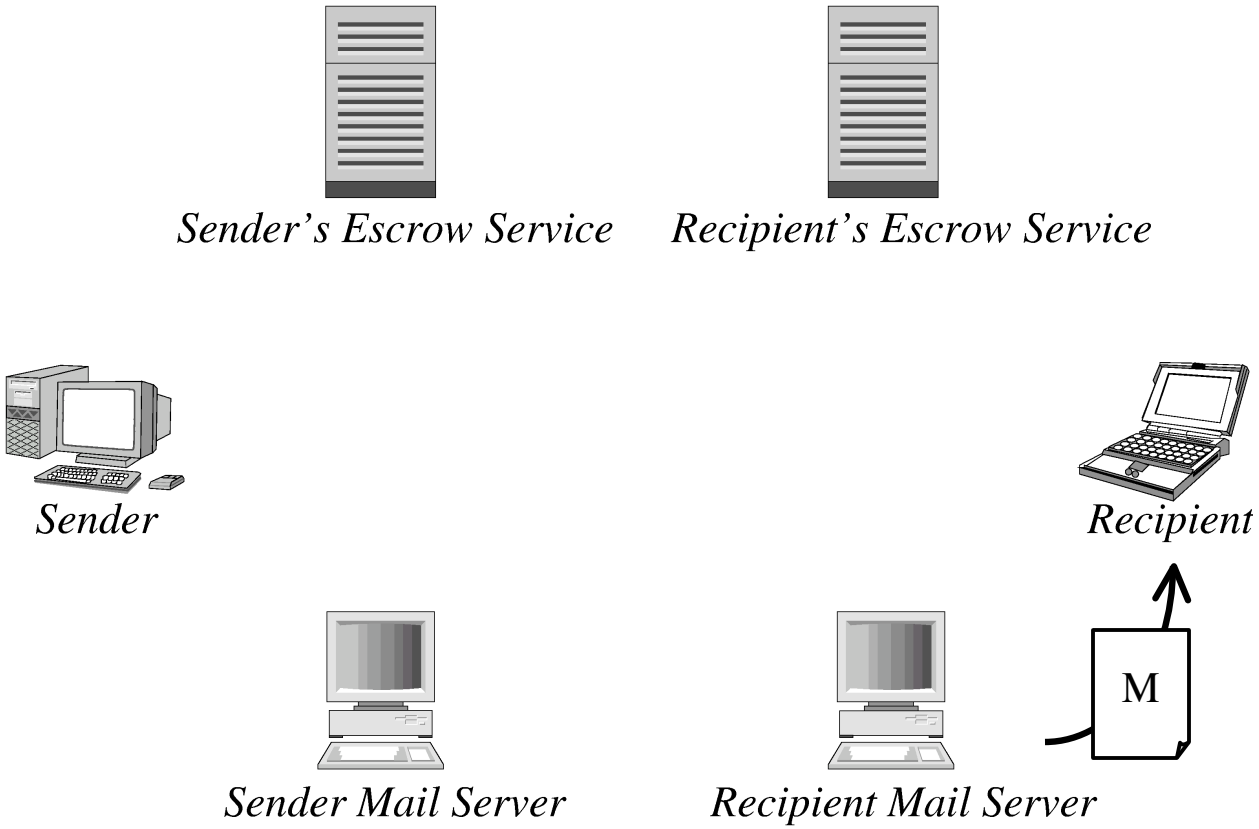

If sender is whitelisted, mail is delivered

The ABM in Action

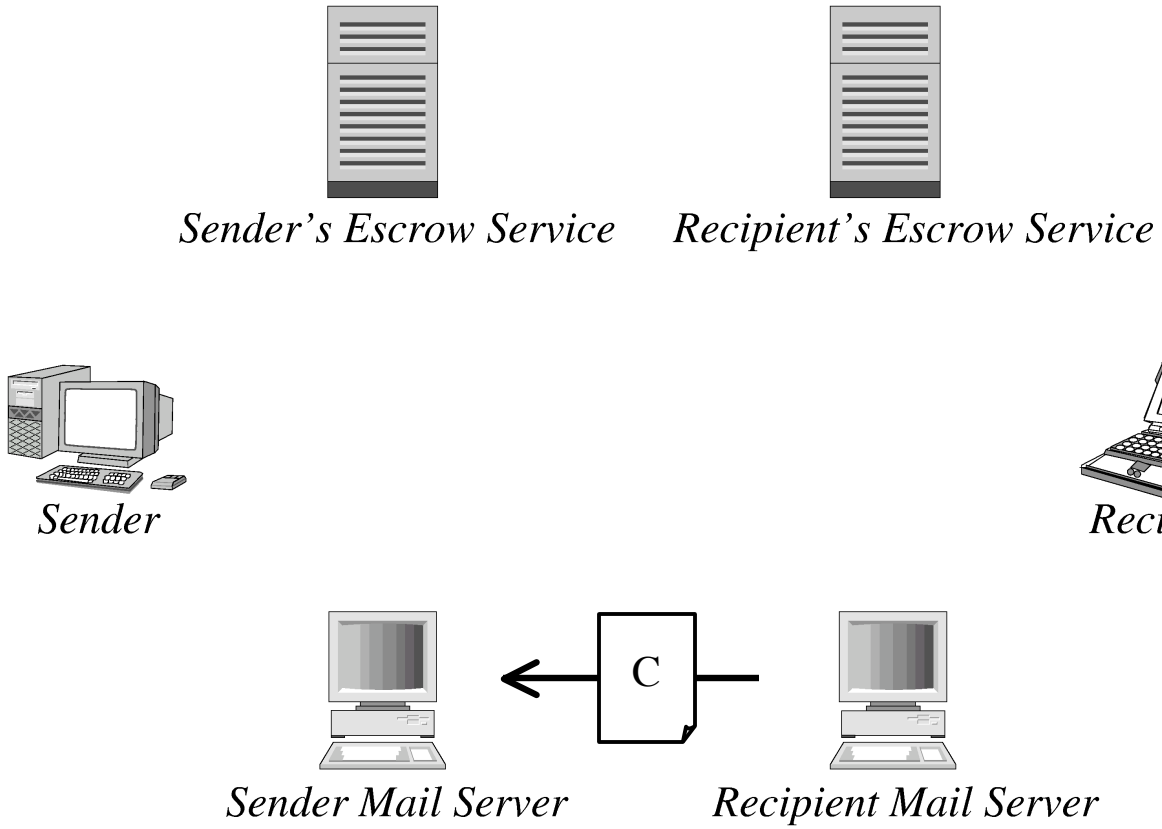

If sender is not whitelisted, a bond-challenge is issued

The ABM in Action

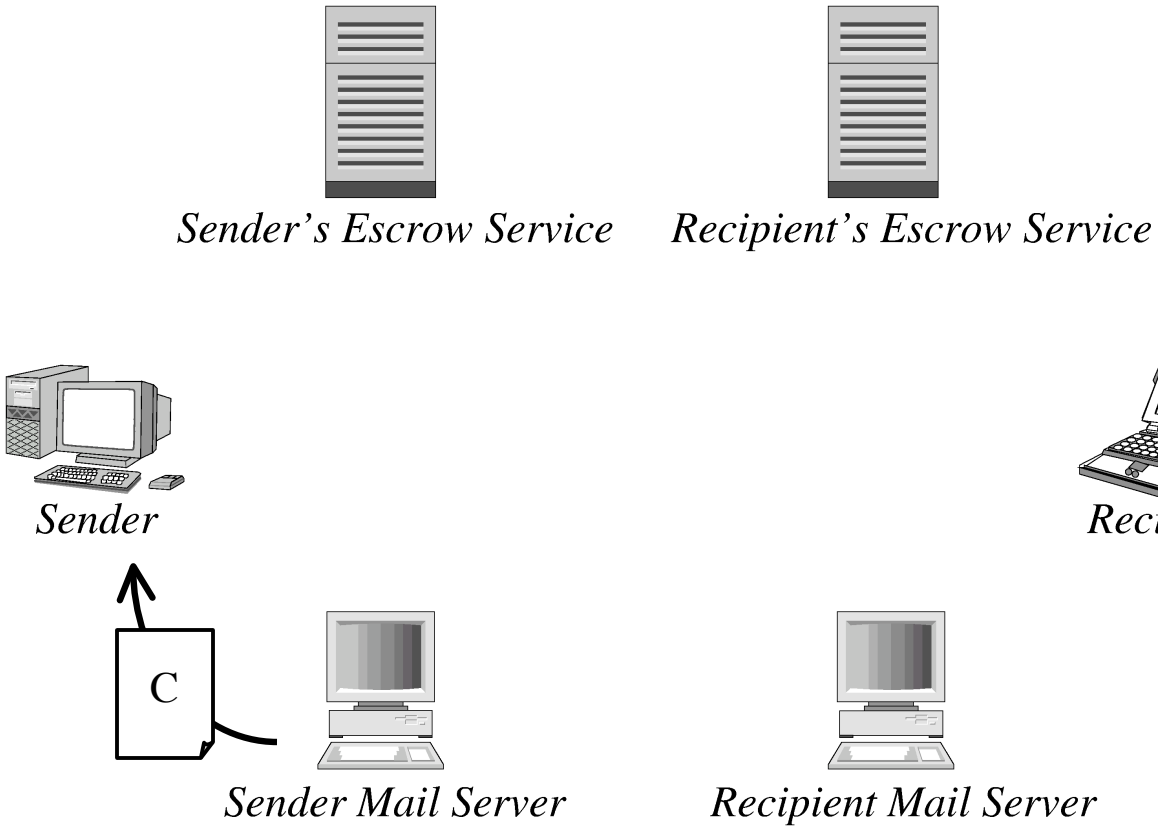

The challenge is returned to the sender…

The ABM in Action

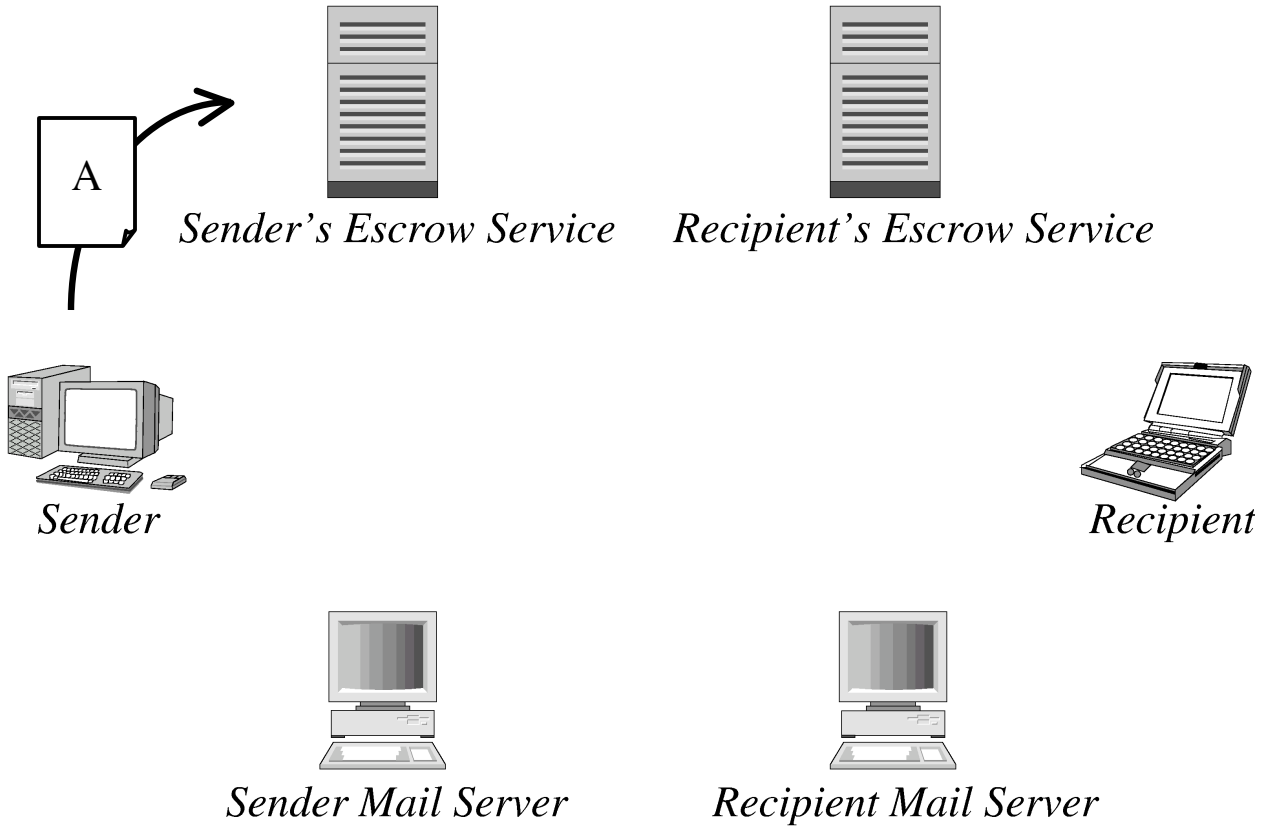

… and the sender can authorize posting a bond

The ABM in Action

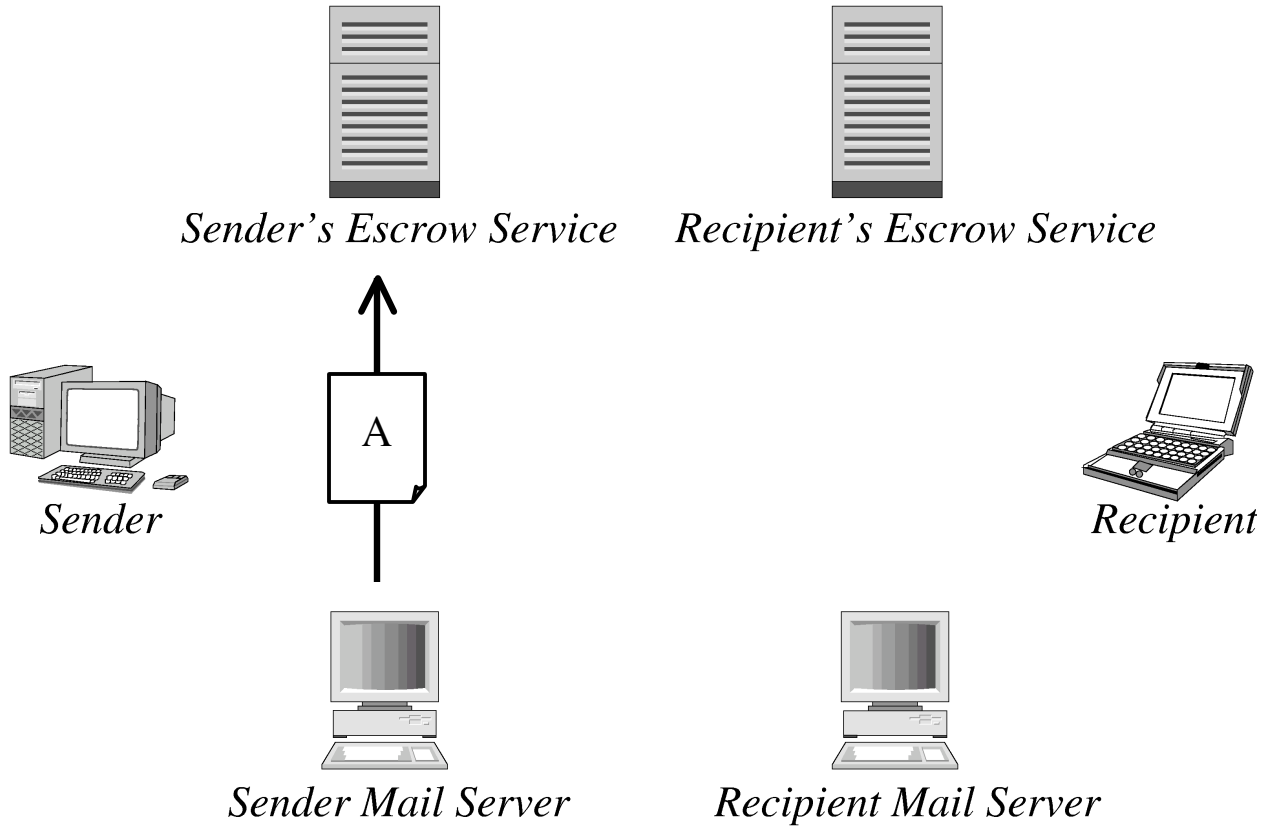

Or, with automation, sender's mail server authorizes bond

The ABM in Action

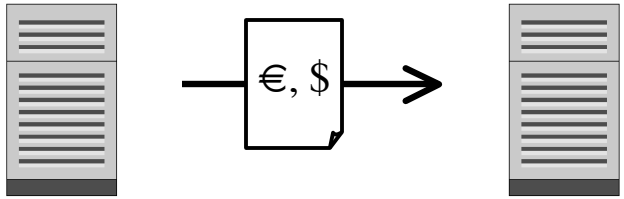

Sender's Escrow Service *Recipient's Escrow Service*

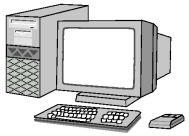

Sender

Recipient

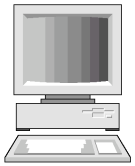

Sender Mail Server

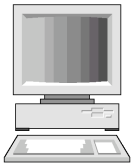

Recipient Mail Server

The bond is posted

The ABM in Action

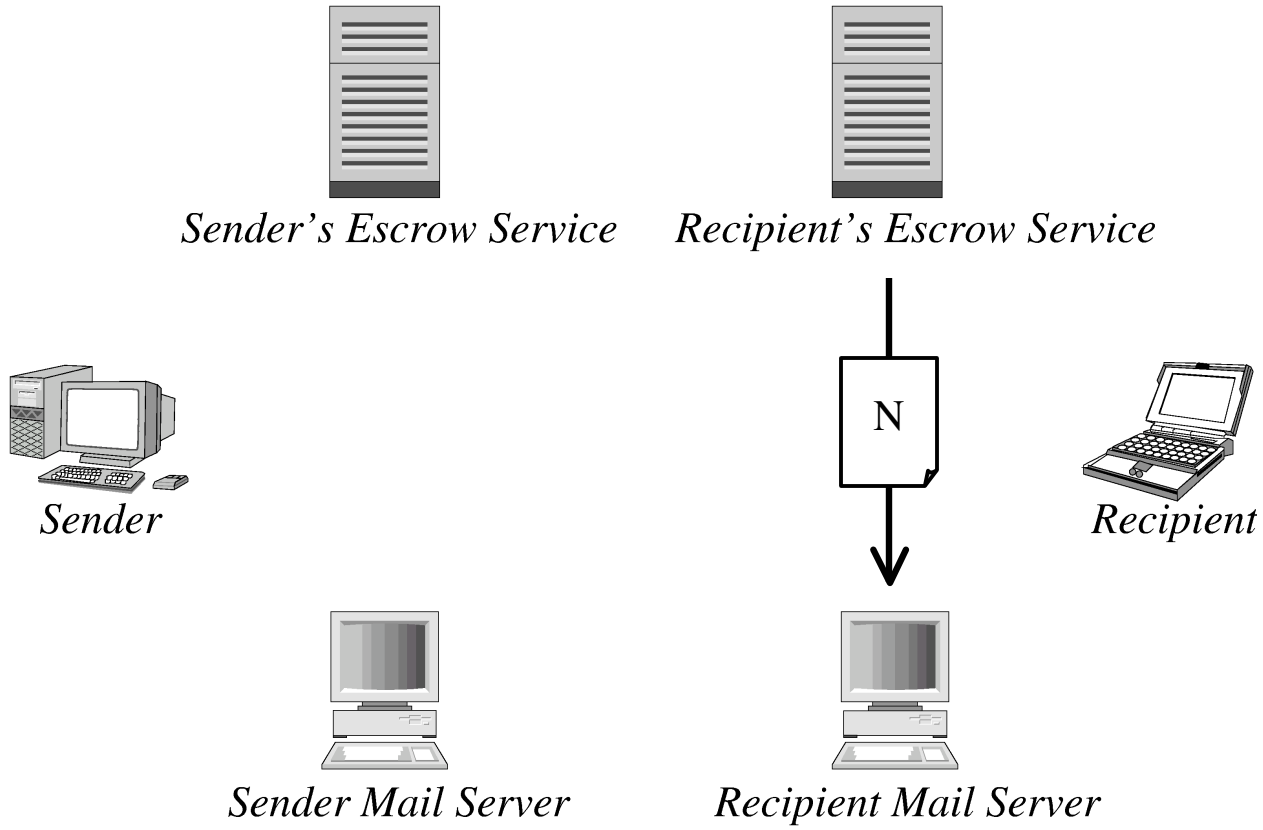

Recipient's Escrow notifies Recipient Mail Server of bond receipt

The ABM in Action

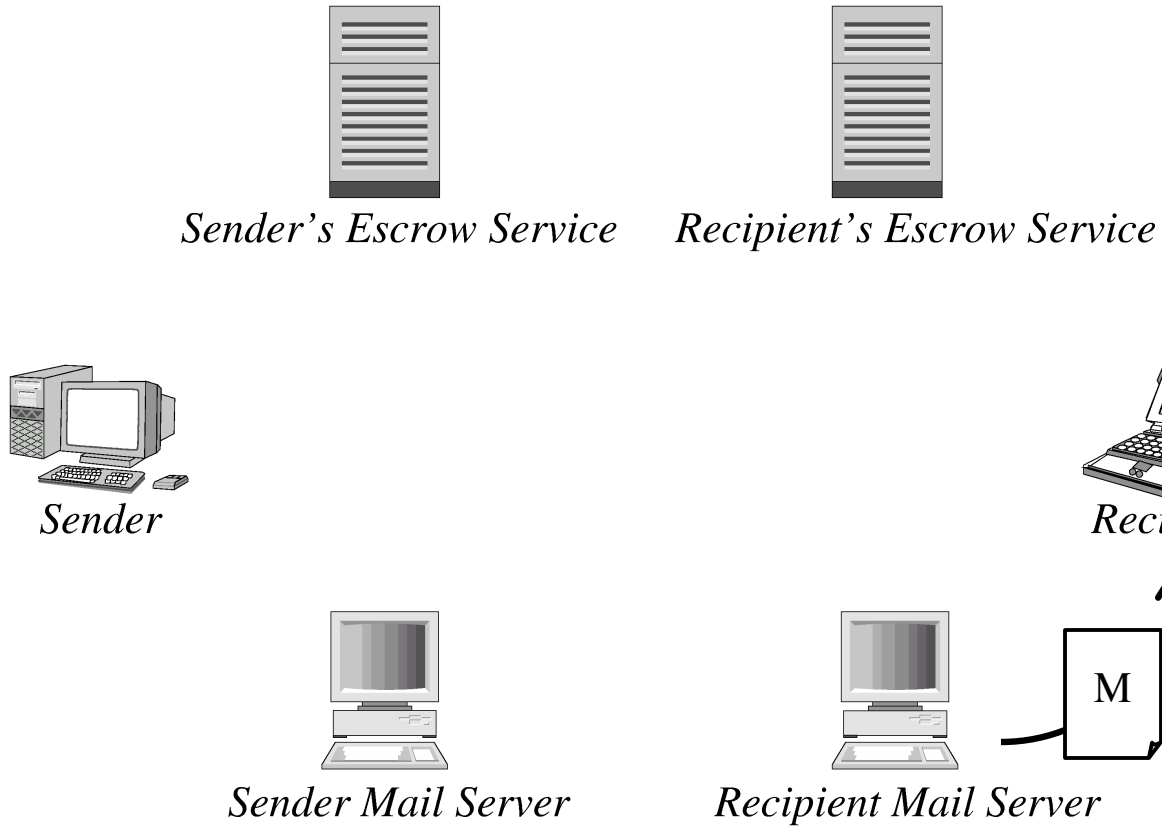

Recipient's Mail Server delivers original message

The ABM in Action

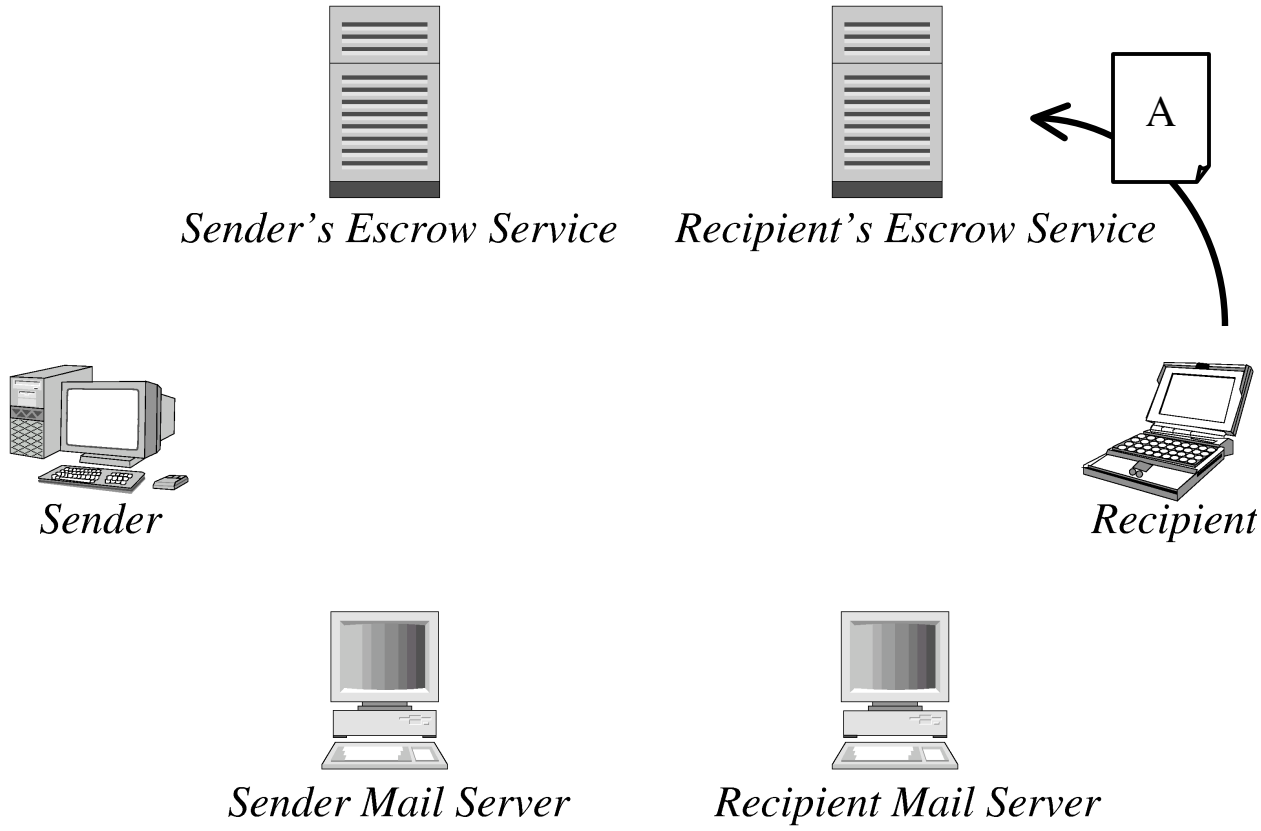

The Recipient claims bond, lets it expire, or…

The ABM in Action

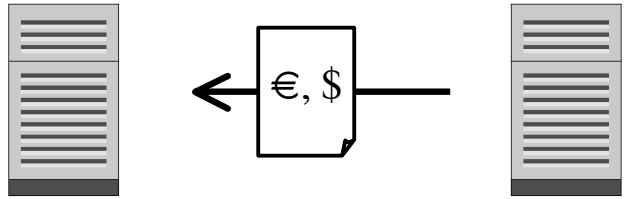

Sender's Escrow Service *Recipient's Escrow Service*

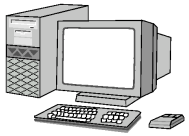

Sender

Recipient

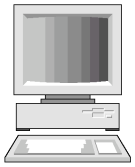

Sender Mail Server

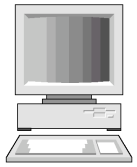

Recipient Mail Server

the bond is returned unclaimed

Value Chain Members

- Senders (Marketers, Organizations, Individuals)
- Escrow Agencies (sender and recipient)
- Underwriters
- Transport (ISPs, Enterprises)
- Recipients

Industry Organization

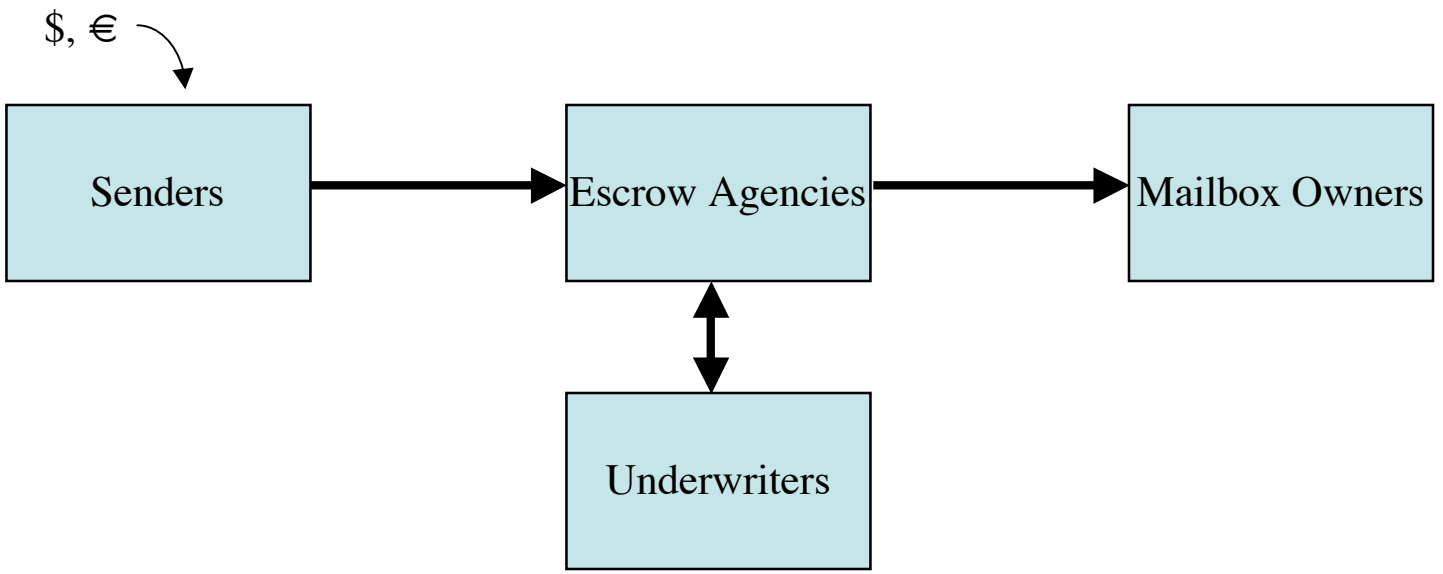

ABM Recap

- Bond is a ‘quality warranty’ collected *before* primary message is exchanged
- A contingent liability with expiration date
- Software-executed ‘policies’ allow cheap negotiation and discovery of mutually favorable terms of exchange

Comparison of Attention Bond to Perfect Filter

- *G* distribution expected to be mostly 'Good'
- *B* distribution expected to be mostly 'Bad'
- Sender does not know *r*
- Sender knows to which distribution (*G* or *B*) his email belongs
- Bond size is the same for both distributions
- Recipient can have seize 'policy'

Filtering Review

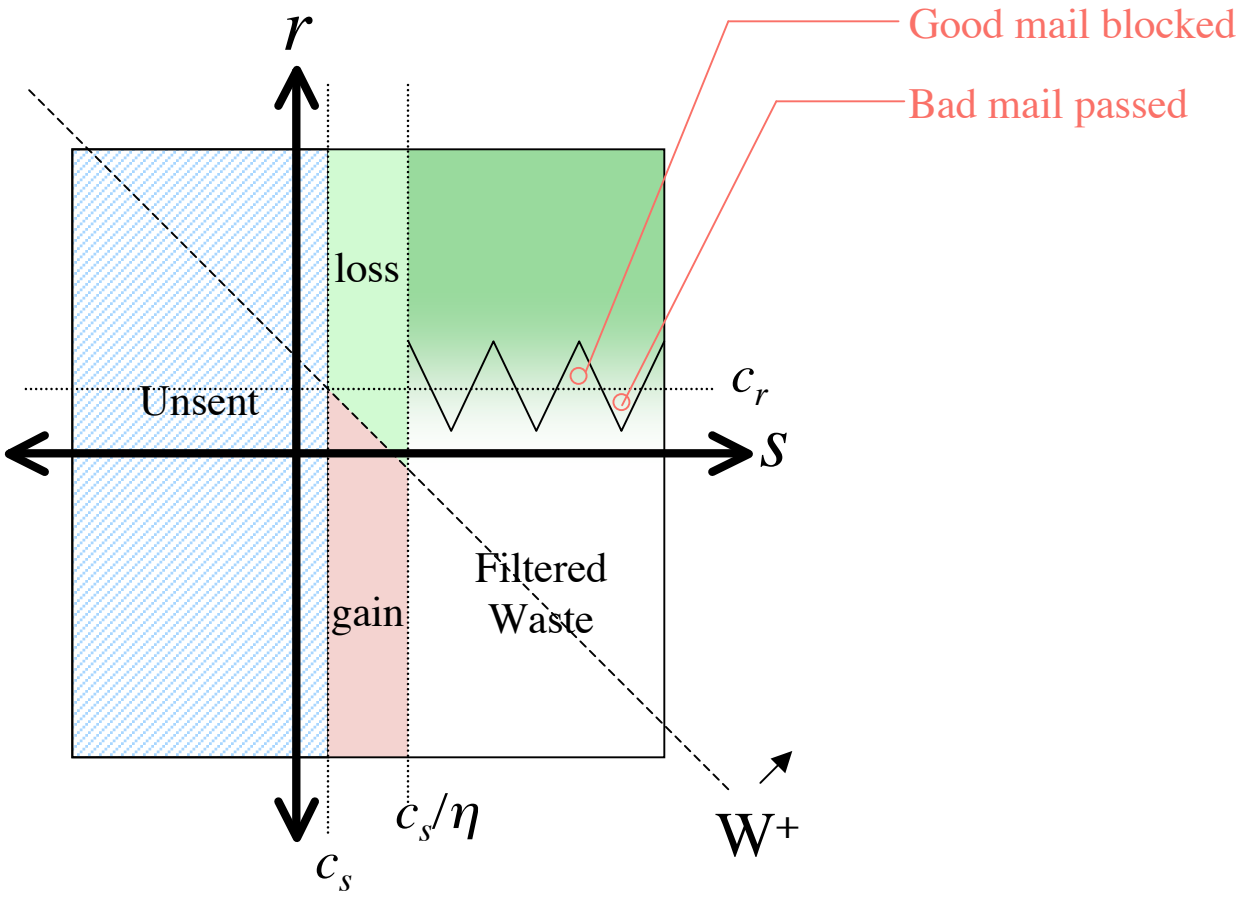

Perfect Filter

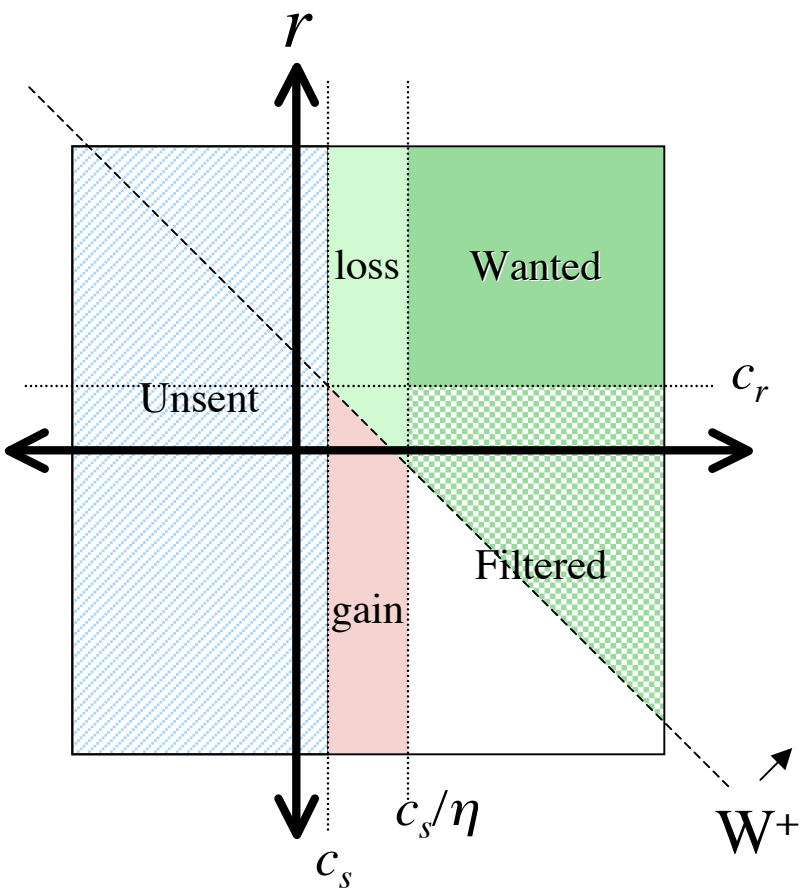

Definition

A technological filter that:

- Operates without cost
- Makes no mistakes (no false positives, no false negatives)
- Intuits and internalizes all reader preferences
- Eliminates, prior to receipt, any email where $r < c_r$

Not all filtered email is waste!

Policies

- Let a policy be any decision rule that results in a specific ex post probability p of seizing the bond.

- Examples
  - Always seize it.
  - Only seize for spam.
  - Only seize for the most offensive spam.
  - Only seize it for your buddies…

- So let the expected $b=p\phi$ where ϕ is the ex ante choice of bond fee and p is the ex post policy on seize rate.

Bond collection policy permits discriminatory costs

- Assume maximum bond size ϕ. If seized with probability p, then expected bond size $b = p\phi$
- Result: $p_G^+\phi < p_B^+\phi$
- More costly for spammers since bond is collected more often

Distributions G and B

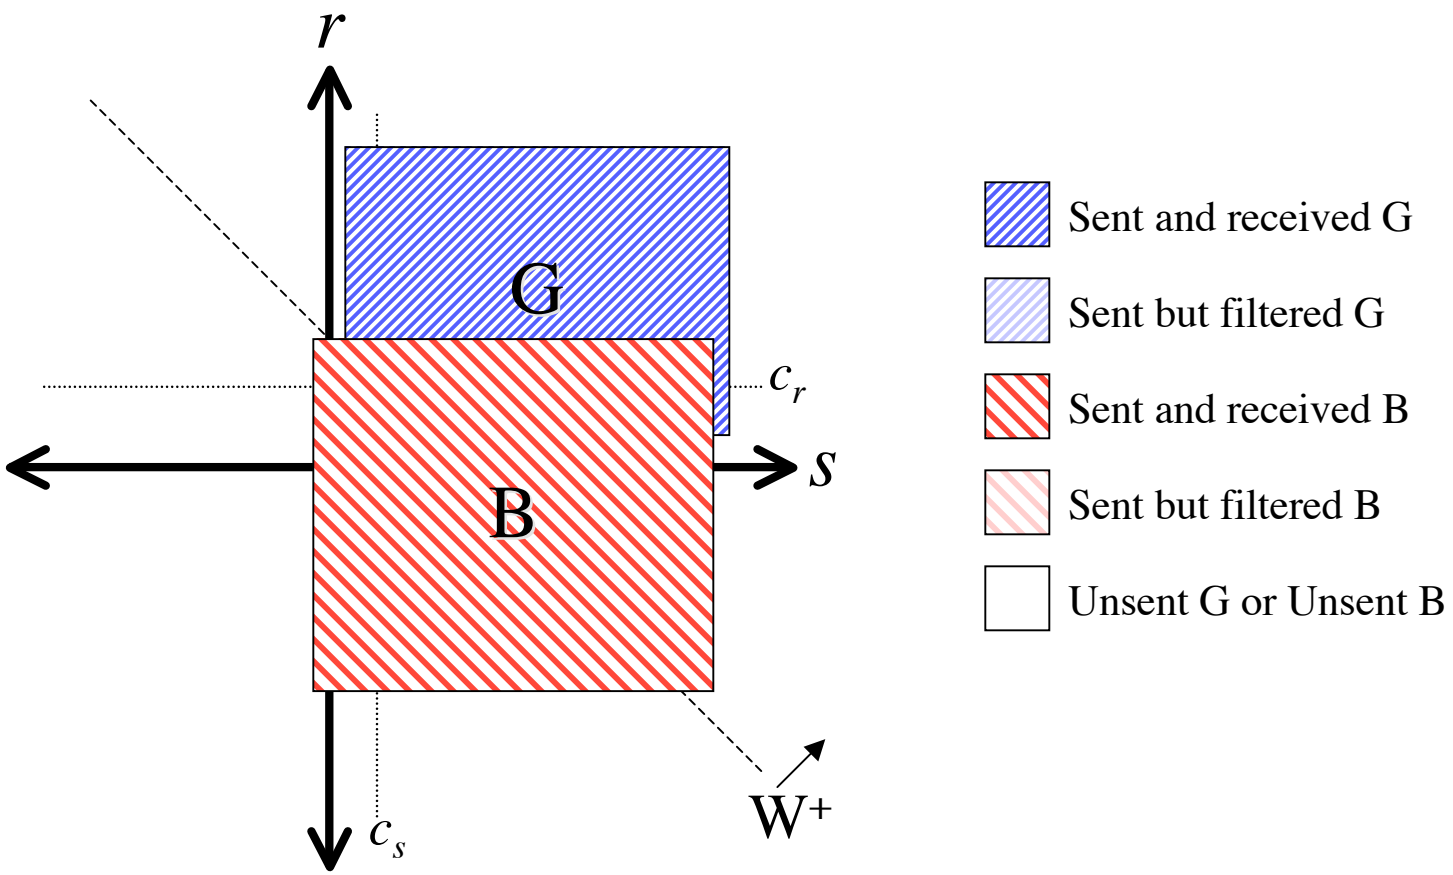

Null Mechanism

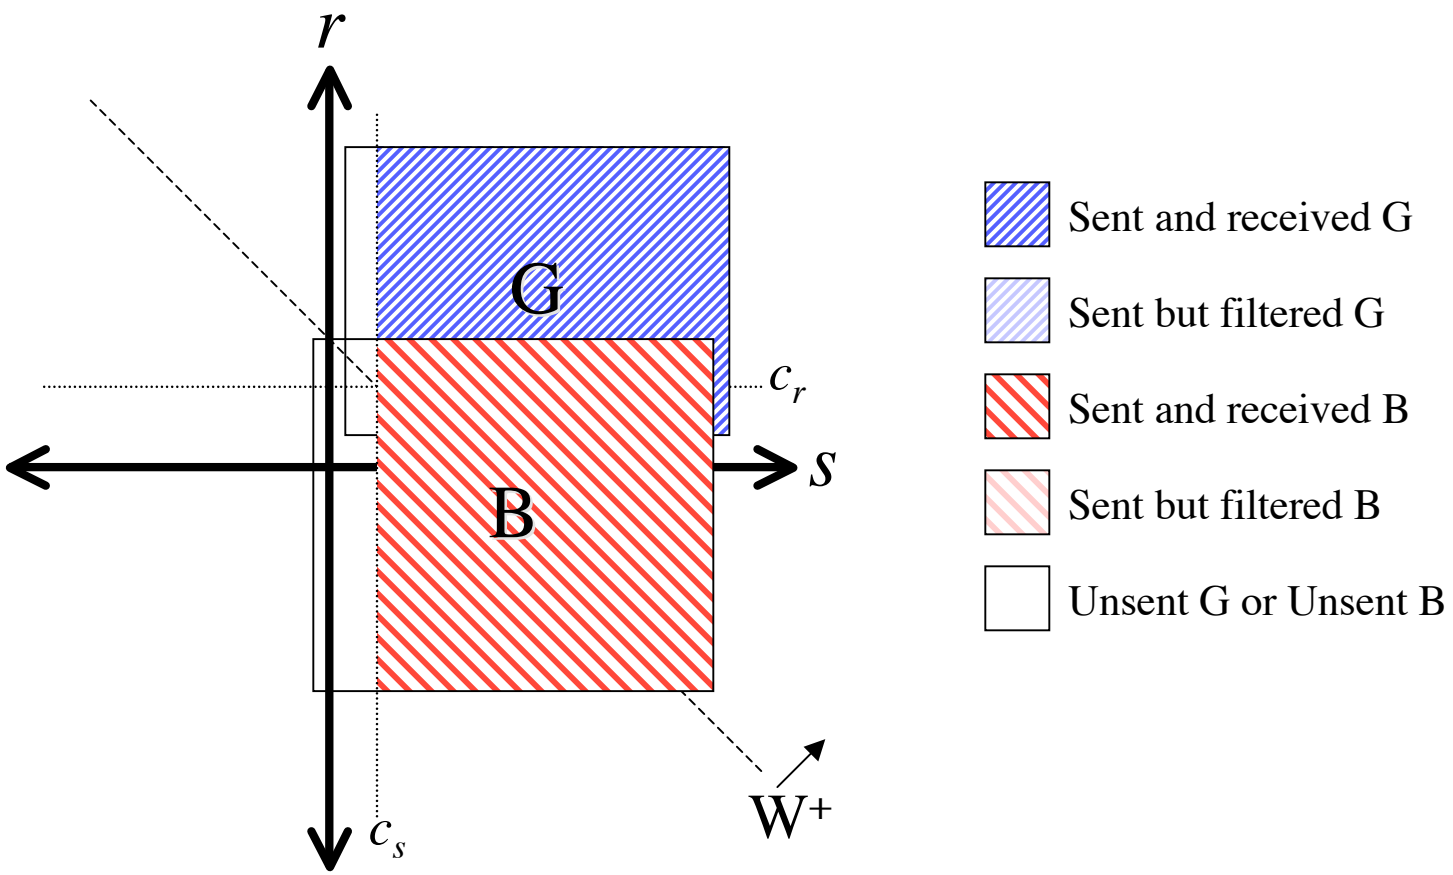

Perfect Filter

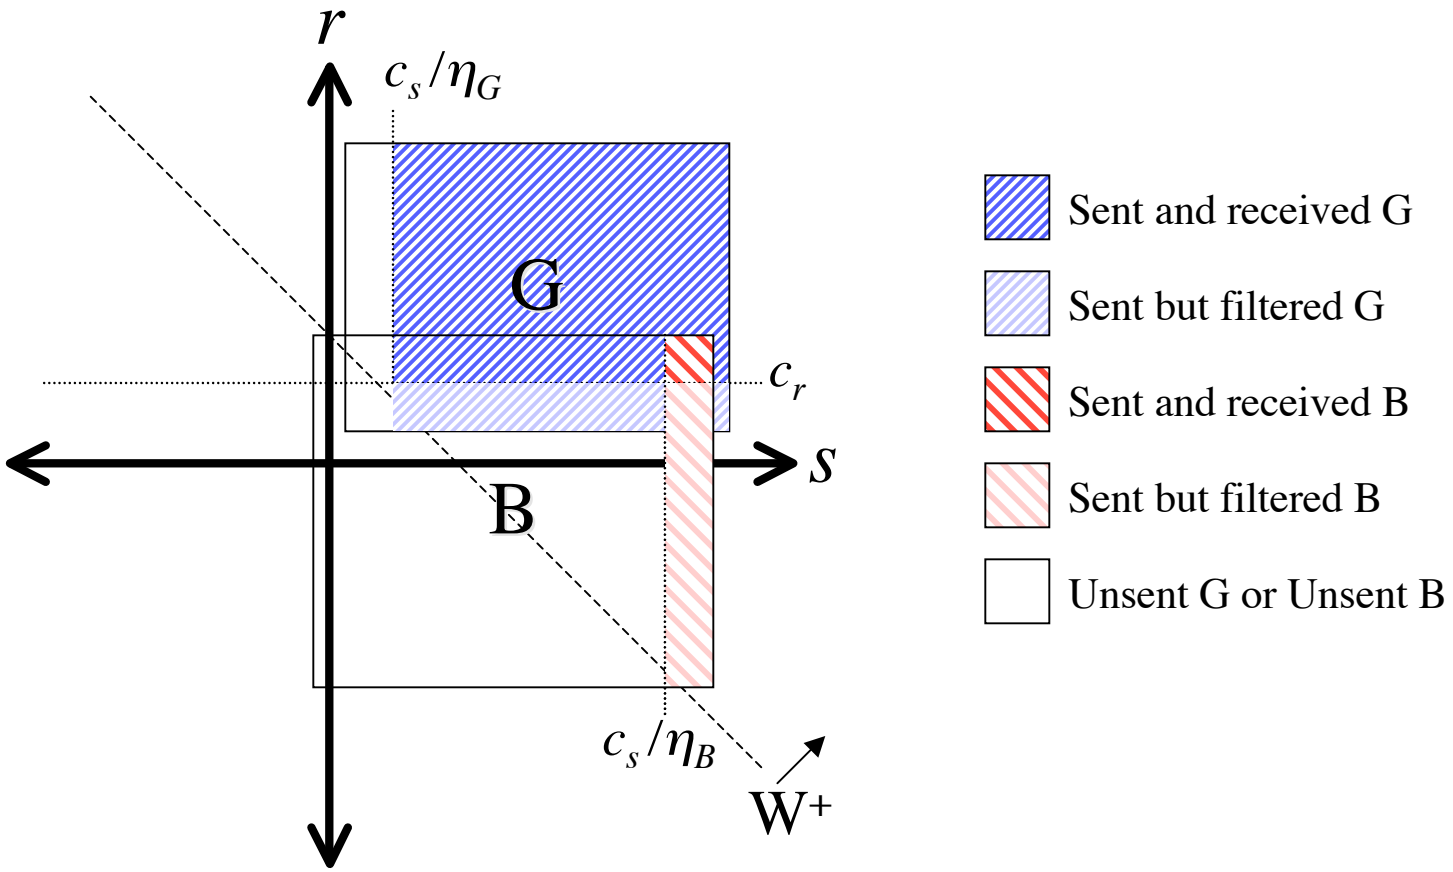

Attention Bond Mechanism

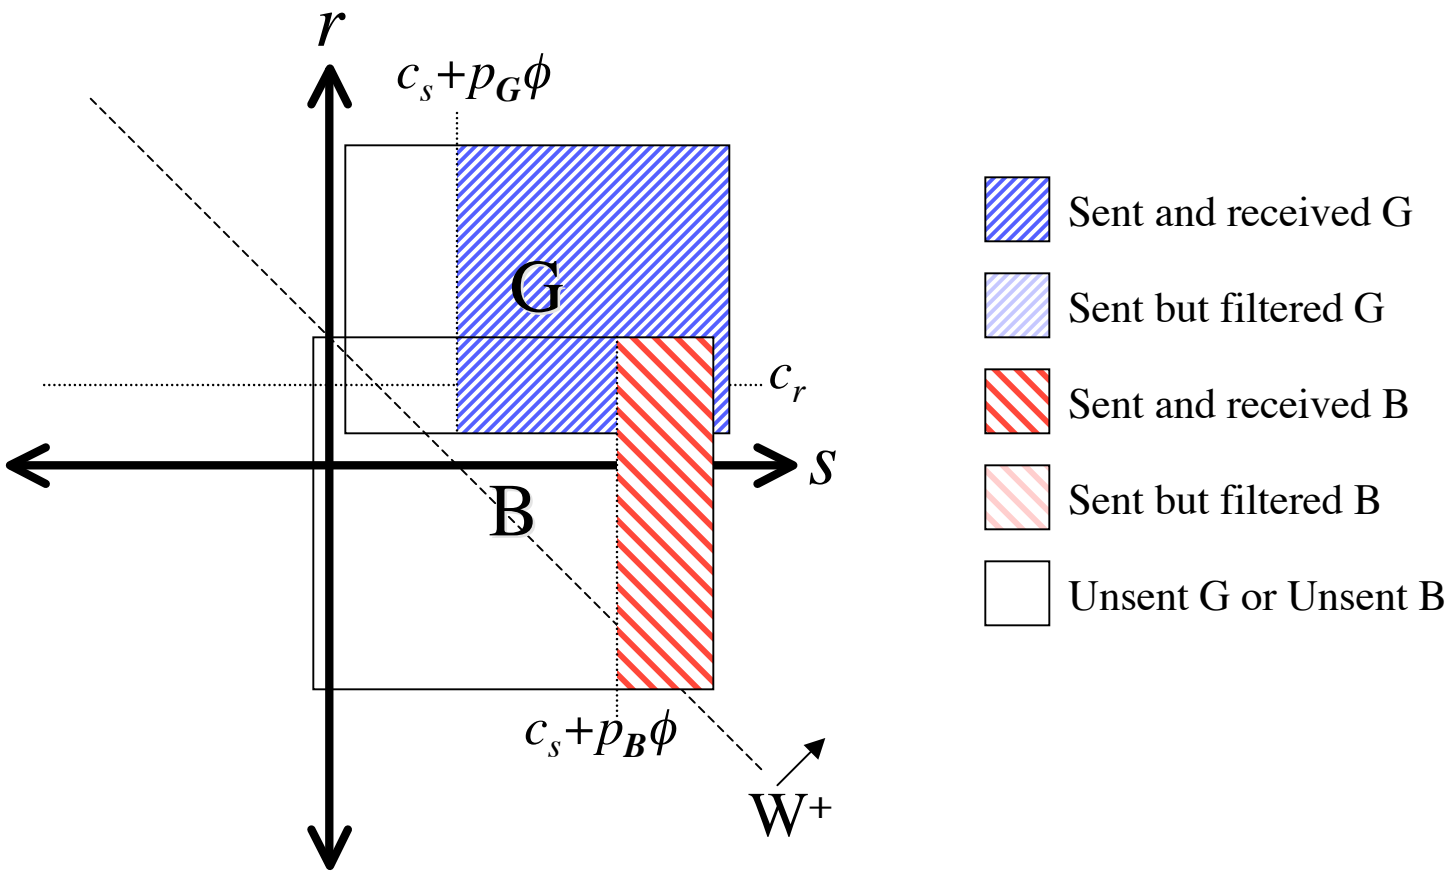

Bond vs. Computation

- Computation is a tax with no benefit to redistribute
- *Transfers* create the possibility of wealth through trade. Transfers need not solely be currency.
- Adoption effects are stronger: adopting mailbox owners are compensated
- Readers can adjust screen to their own private *value* - why should the value of your wasted time depend on someone else's costs?
- All in-bound mail treated equally with ex-ante computation, but bonds promote higher penalties for the bad guys.
- Computation is more susceptible to viral attack: you may be able to steal CPU cycles without the victim noticing (like they do now!) but not money from a capped account.
- Easily "signals" really important messages.

What about e-stamps?

A stamp can look like an Attention Bond if

- The recipient gets the proceeds (not ISP or gov't)
- The recipient chooses how much the postage costs
- The recipient can costlessly refund the postage
- A recipient controlled whitelist exempts familiar senders

Why not just use a bond?

ABM and Direct Marketing

- Cheaper than traditional mail channels
- Recipient gets bulk of benefit, not printer and mailer
- Greater information feedback than TV, Radio, Snail Mail, or Magazines
- Recipient gets off mailing lists easily by claiming bond

Caveats and Adoption Issues

- Requires Infrastructure
  - Escrow service(s), payment network, underwriters
  - Protocol enhancements (on top of SMTP/S-MIME)
  - Server infrastructure, Whitelist management
  - Authentication, strong or “stronger” identities
- Transaction Costs?
  - Debatable. We have phone systems that bill to the second
- Network Effects
  - Market fracture could slow adoption
  - Standards a must

How might adoption happen?

- Creation of standards (leveraging others - S/MIME, x509, OpenPGP, Caller-Id, transition to stronger authentication)
- Establishment of financial intermediaries (e.g. PayPal, credit cards, banks, ISPs, Hotmail, Google, Yahoo)
- Consumer/Provider client & server rollout

Policy & Enforcement Issues

- Need consumer fraud protection analogous to that for Credit Cards
- Need escrow account clearing analogous to Automated Clearing House (ACH) for checks
- Need pro-competitive open standards
- Need privacy regulations for transaction histories

Long-term Social Benefits

- Boon to commerce: replace threat of communications veto with option for fruitful exchange
- Maintain low cost nature of medium, provide individual control
- Direct marketing
  - Cheaper than traditional mail channels
  - Receiver gets the benefit, not printer and mailer
  - Costs less than snail-mail
- Reduce Friction: Recipients have reason to publish contact information, not hide it
- Shifts arms race to crypto.
- Tailors to individual preferences
- New opportunities for industry

Conclusion

- Screening action of bond forces senders to reveal their intentions
- Creates possibility of greater trade volume, with both parties wanting to participate
- Facilitates wealth transfers, can benefit recipients more than unilateral veto.
- *Can beat a Perfect Filter.*
- Individualized, works with existing uses
- Returns control of the mailbox to its owner

Next Steps

- Open standards for:
  - Challenge message, notification of payment formats
  - Inter-escrow agency payments/transfers
- Establishment of first Consumer and Marketer-focused escrow agencies
- 'De-regulated' style regulation of agencies, creation of market for relevant services
- Deployment of protocol implementing servers and clients
- Whitelist management tools sets (individual, enterprise/organization)

Questions?

For further info see:

1. "An Economic Answer to Unsolicited Communication", Proceedings of the 5th ACM conference on Electronic Commerce, 2004. Loder, Van Alstyne, Wash.
2. FAQ and Links: http://www.eecs.umich.edu/~tloder/abm_faq.html

Alan Ralsky: "There is too much money involved. I'm a survivor. And when you are a survivor, you find a way to make it happen."

Speaking to NYT – 12/30/03

Larry Lessig: "Congress's efforts to protect against spam will have failed"

Making predictions on NPR for 2004

Timothy Muris: "I'd advise customers not to waste their time and effort. Most spam is already so clearly illegitimate that the senders are no more likely to comply with new regulations than with the laws they now ignore."

Speaking on DO NOT SPAM lists

The Reader's choice of bond

Reader Surplus defined: $$RS_B = k\int_{\underline{r}}^{\bar{r}}\int_{c_s+b}^{\bar{s}}(r - c_r + b)\,dr\,ds$$

Optimal expected bond: $$b^+ = \frac{1}{2}\left((\bar{s} - c_s) - \left(\frac{\bar{r} + \underline{r}}{2} - c_r\right)\right)$$

Reader Surplus: $$RS_B = \frac{1}{4(\bar{s} - \underline{s})}\left((\bar{s} - c_s) + \left(\frac{\bar{r} + \underline{r}}{2} - c_r\right)\right)^2$$

Always: $$RS_B \geq RS_0$$

The Perfect Filter

Perfect Filter: $\eta = \dfrac{\bar{r} - c_r}{\bar{r} - \underline{r}}$ Send if: $\eta s \geq c_s$

Reader Surplus defined: $RS_{PF} = \dfrac{k}{\eta}\int_{c_r}^{\bar{r}}\int_{c_s}^{\eta\bar{s}} (r - c_r)\,dr\,ds$

Reader Surplus: $RS_{PF} = \dfrac{(\bar{r} - c_r)^2(\eta\bar{s} - c_s)}{2\eta(\bar{v}_r - \underline{v}_r)(\bar{v}_s - \underline{v}_s)}$

$RS_B \geq RS_{PF}$ Bonding wins if: $b^2 \geq 2(\bar{v}_s - c_s)\left(\dfrac{\bar{v}_r + \underline{v}_r}{2} - c_r\right)$