The Welfare Consequences of Fake Reviews

Ashvin Gandhi (UCLA) and Brett Hollenbeck (UCLA)

FTC, November 2023

Motivation

Online Reputation Systems:

- Large literature showing significant causal effect of seller ratings on sales on many platforms
- Strong incentives for sellers to manipulate their reputations
- Two-sided platform structure makes regulation difficult

Consequently, rating manipulation is common in e-commerce, arguably worse than ever.

Motivation

Online reputation systems:

- Valuable mechanism for solving asymmetric information problem in online markets and platforms (Tadelis (2016))
- Ratings and reviews benefit both sellers and consumers (Reimers & Waldfogel (2021))

Regulators increasingly see rating manipulation as an important and growing problem for consumer protection:

- FTC - Proposed rule in June 2023 currently
- UK CMA - Proposed law regarding fake reviews
- EU Digital Services Act (DSA) - stricter new regulations on fake reviews

This Paper

Question: What are the impact of fake reviews on Amazon.com?

Approach:

1. Provide a *framework* for assessing the impact of fake reviews.
2. Gather *data* on fake reviews and on consumers' perceptions.
   - Real products actually purchasing of fake reviews
   - Incentivized elicitation of consumers' beliefs
3. *Estimate* an empirical model
   - Ratings, demand, pricing
4. *Simulate* counterfactual enforcement against fake reviews.

Outline of Talk

1. **Simple model of welfare**
2. Data and setting
3. Empirical model of beliefs from ratings and demand
4. Counterfactual Results

Channels for Effects

1. **Misinformation**
   - Fake reviews mislead consumers into purchasing undesirable products.
   - Fake review purchasers (FRPs) can raise prices
   - Honest products (NFRPs) must lower prices

Channels for Effects

1. **Misinformation**
   - Fake reviews mislead consumers into purchasing undesirable products.
   - Fake review purchasers (FRPs) can raise prices
   - Honest products (NFRPs) must lower prices

2. **Mistrust**
   - Erodes long-term trust in ratings.
     - Only 17% fully trusts reviews
   - No longer solves the asymmetric information problem.
     - Lower demand and potentially less sensitive to ratings
   - Greater overall price competition

Fake Review Purchaser & Unaware Consumer

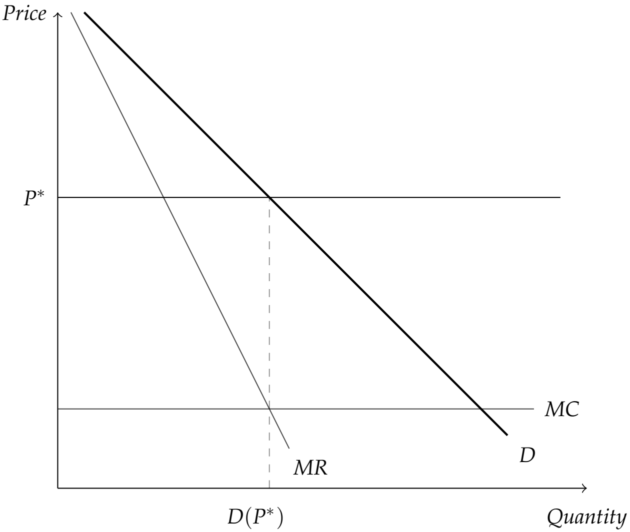

Econ 101: optimal price given a demand curve and marginal cost.

Fake Review Purchaser & Unaware Consumer

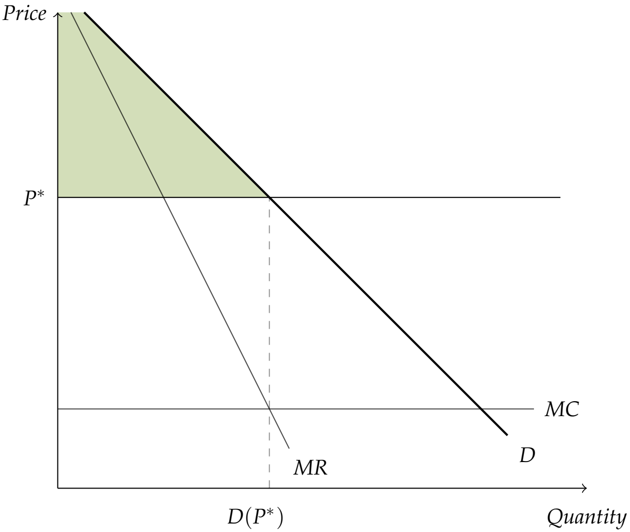

Econ 101: consumer surplus.

Fake Review Purchaser & Unaware Consumer

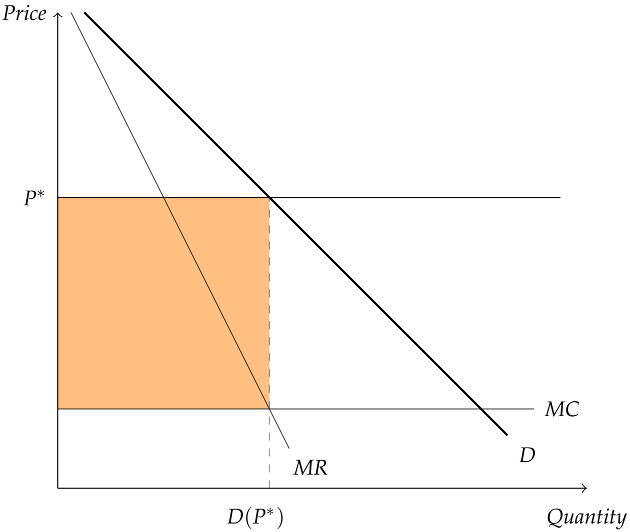

Econ 101: producer surplus.

Fake Review Purchaser & Unaware Consumer

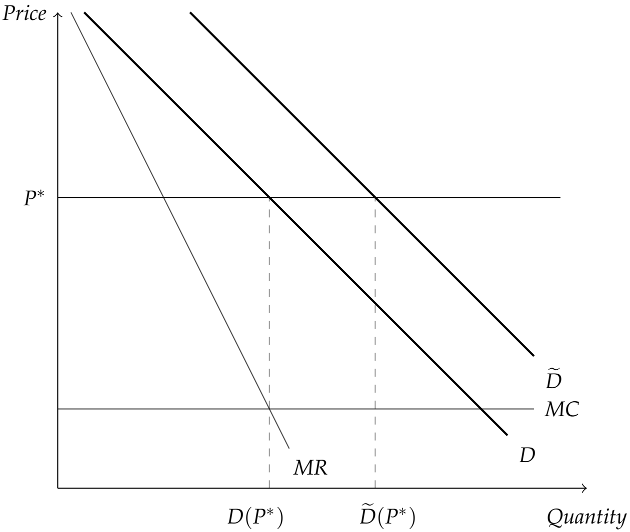

But what if the firm purchases fake reviews to face demand $\tilde{D}$?

Fake Review Purchaser & Unaware Consumer

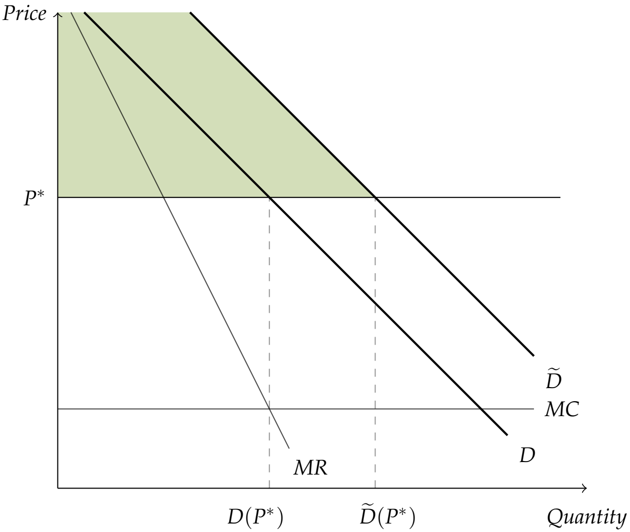

Misinformation: *Holding price fixed*, consumers believe they'll get this.

Fake Review Purchaser & Unaware Consumer

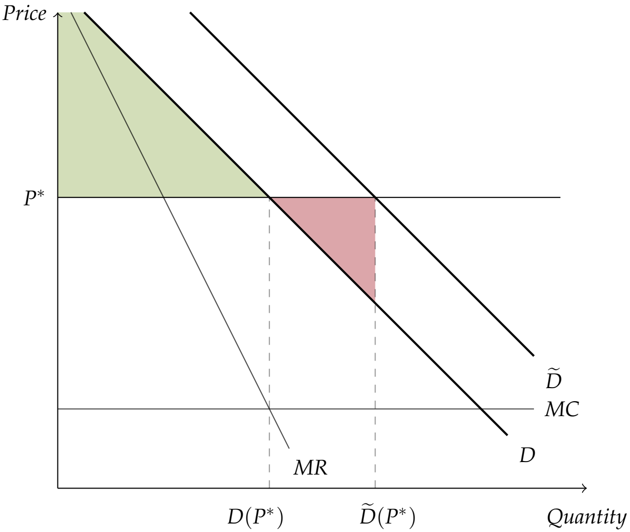

Misinformation: Instead they actually receive this.

Fake Review Purchaser & Unaware Consumer

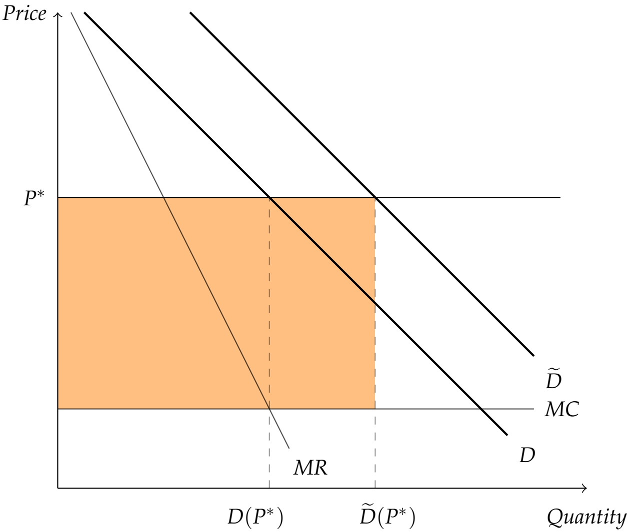

Misinformation: The firm's profit increases.

Fake Review Purchaser & Unaware Consumer

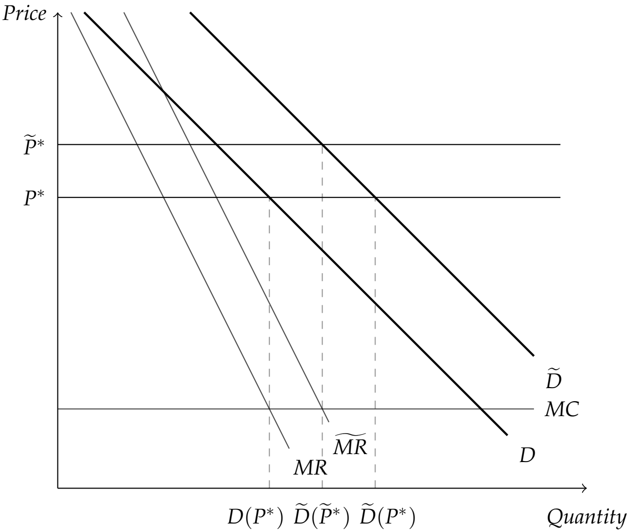

Pricing: But the firm should increase its price.

Fake Review Purchaser & Unaware Consumer

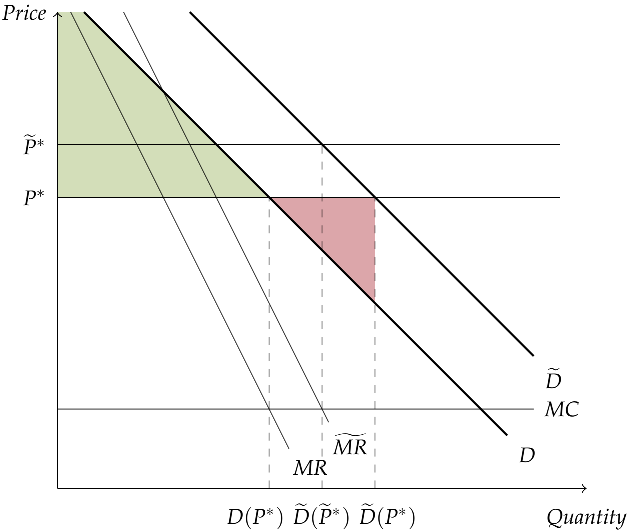

Pricing: Recall consumer surplus before the price increase.

Fake Review Purchaser & Unaware Consumer

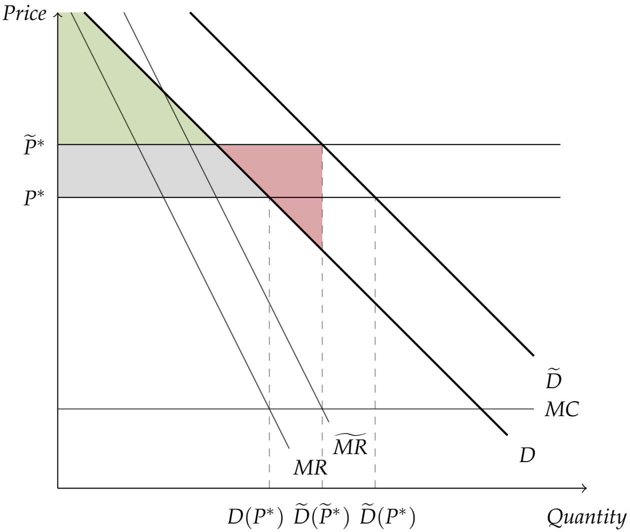

Pricing: The price increase further lowers consumer surplus.

Honest Product & Unaware Consumer

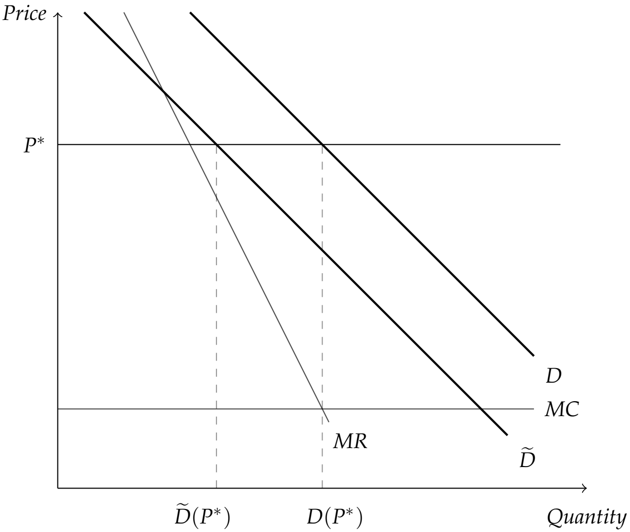

But what if the *firm's competitor* purchases fake reviews?

Honest Product & Unaware Consumer

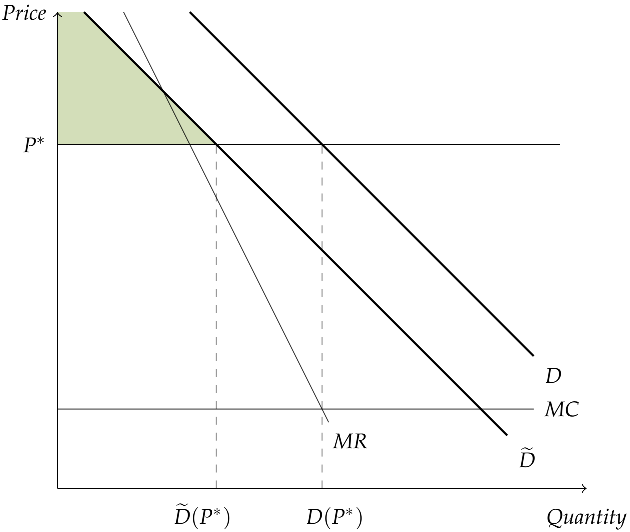

Misinformation: *Holding price fixed*, consumers believe they'll get this.

Honest Product & Unaware Consumer

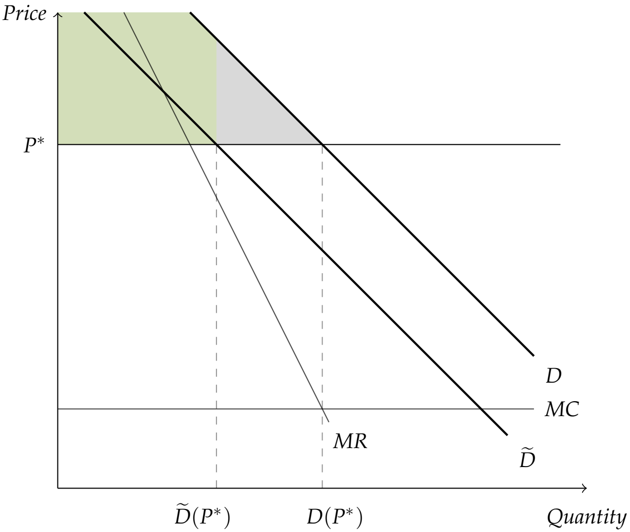

Misinformation: Those that do buy are better off than they knew.

Honest Product & Unaware Consumer

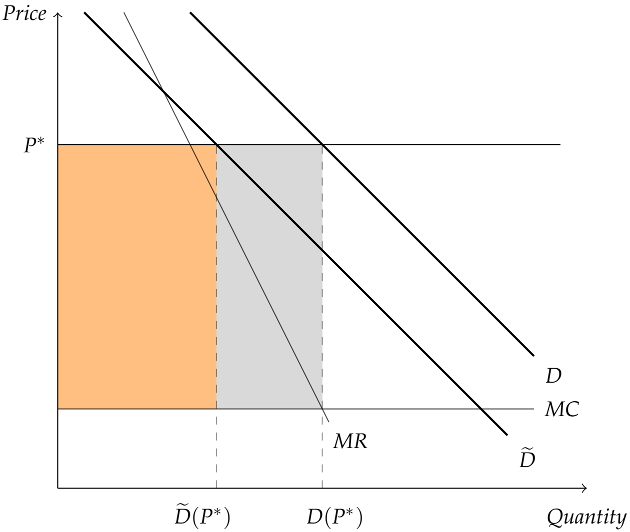

Misinformation: The firm's profit decreases.

Honest Product & Unaware Consumer

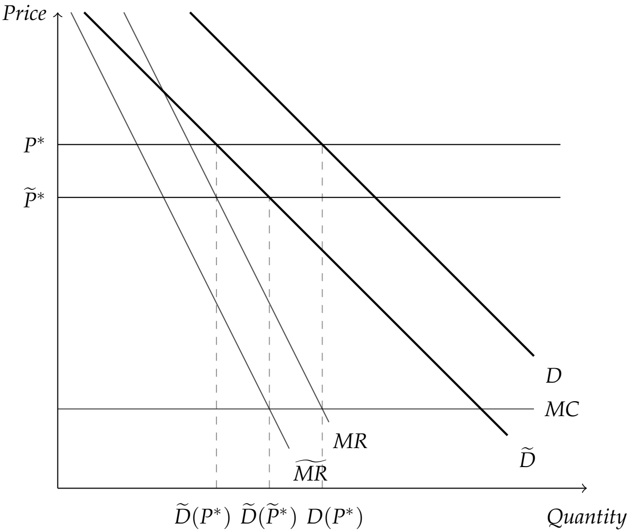

Pricing: But the firm should decrease its price.

Honest Product & Unaware Consumer

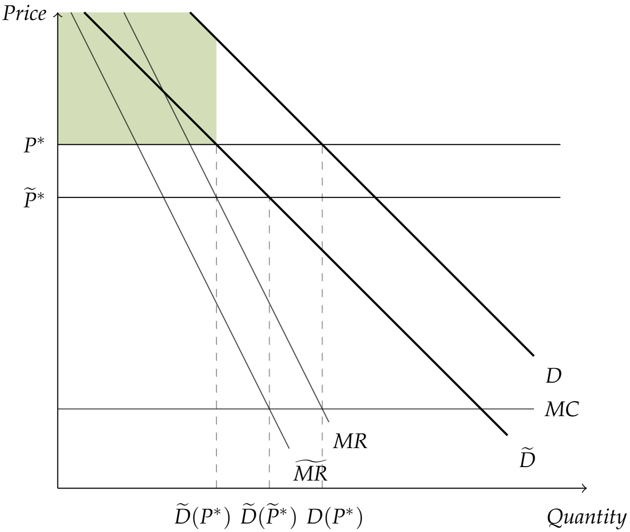

Pricing: Recall consumer surplus before the price decrease.

Honest Product & Unaware Consumer

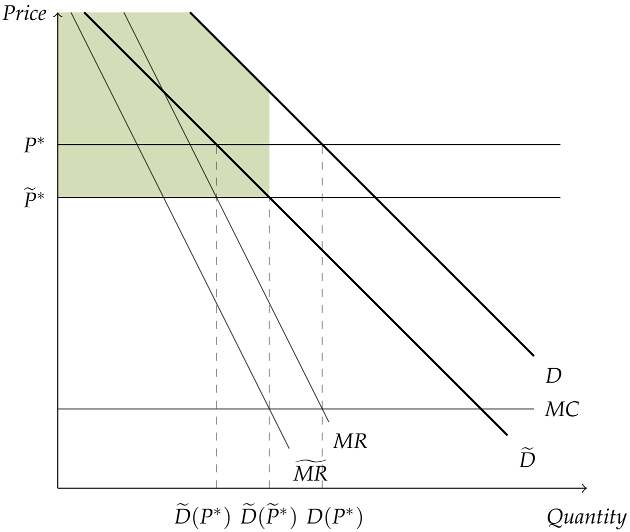

Pricing: The price decrease increases surplus. (Above no fake reviews!)

Awareness of Manipulation Causes Mistrust of Ratings

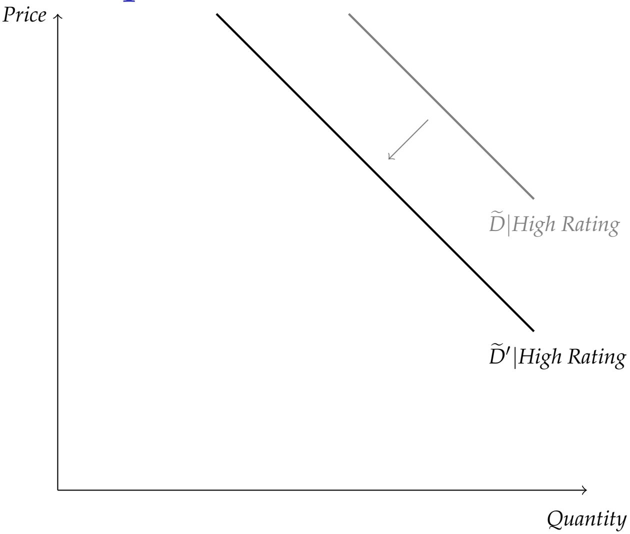

Mistrust: A high rating might simply reflect fake reviews.

Awareness of Manipulation Causes Mistrust of Ratings

Mistrust: A competitor's high rating could reflect fake reviews.

Awareness of Manipulation Causes Mistrust of Ratings

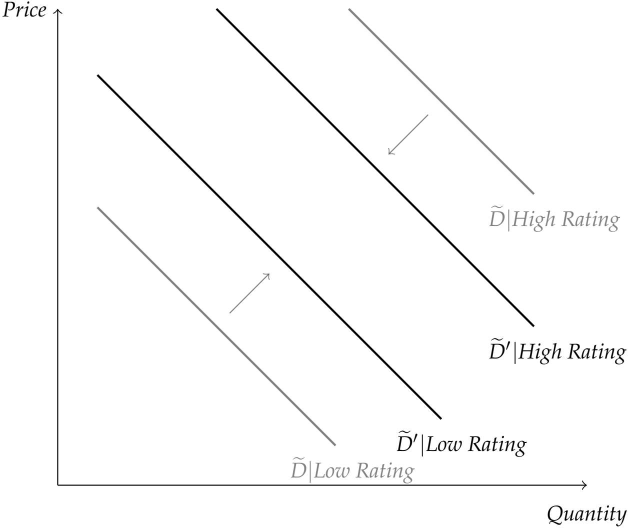

Mistrust: Changes the relationship between ratings and perceived quality.

Outline of Talk

1. Simple model of welfare
2. **Data and setting**
3. Empirical model of beliefs from ratings and demand
4. Counterfactual Results

Markets for Fake Reviews

- Facebook is the largest channel for purchasing fake reviews
  - Large private groups connecting Amazon sellers and reviewers

Process:

1. Amazon seller (or broker) posts brief description and image of product
2. Reviewer responds privately with proof of account, address, etc.
3. Reviewer purchases product and shows proof of a positive review
4. Amazon seller pays reviewer: price + taxes + fees [+ commission]

Markets for Fake Reviews

Need reviews, USA only
PP refund cover pp fees
Pm me if you interested

FREE
shower shelf / toilet paper holder
China
Message

Cordless vacuum
FREE FREE FREE
US only
Half r*fund after shipment and half after r*view
Promise r*fund within 48hrs
if interested，pls PM me

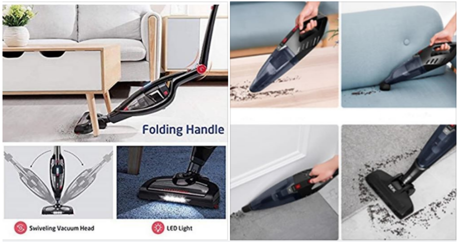

FREE
Cordless vacuum
Message

1
6 Comments

Fake v.s. Incentivized Reviews

Incentivized reviews:

- Incentive is disclosed in the review
  - "Vine Customer Review of Free Product"
- Negative and positive reviews receive equal payment
- Sellers can't pick incentivized reviewers
- $\implies$ payoff is higher if high quality (Li, Tadelis, and Zhao, 2020)

Fake v.s. Incentivized Reviews

Incentivized reviews:

- Incentive is disclosed in the review
  - "Vine Customer Review of Free Product"
- Negative and positive reviews receive equal payment
- Sellers can't pick incentivized reviewers
- $\implies$ payoff is higher if high quality (Li, Tadelis, and Zhao, 2020)

Fake reviews (focus of our paper):

- Payment is not disclosed in review
- Seller requires a 5 star review for reimbursement
- Reviewer often does not test or even open product before reviewing

FB Groups

- Over 4 months we observe top active groups for buying fake reviews
  - 23 groups per week on average
- Average 16,000 members per group
- Average 568 seller posts per group per day
- Crude calculation suggests up to 4.5 million distinct products in a year

FB Data Collection

Group of UCLA undergraduates infiltrate these FB groups and select a random sample of 1400 posts:

1. Identify product on Amazon
2. Collect data on product page and attributes
3. Collect data on other posts by same seller earlier in time
4. Continually search seller/product to find more posts
    - Identify both **start and end date** of fake review recruiting

Amazon Data Collection

Large-scale daily scraping of Amazon.com data on product outcomes:

1. Category-level data from searching product keywords (price, rating, reviews, keyword organic ranks, sponsored listings)
2. Review data for fake review products and close competitors
3. Reviewer data for these products
4. Sales rank (quantity) for all products
   - Chevalier and Goolsbee (2003); He and Hollenbeck (2020)

Who Engages in Rating Manipulation?

Category	N	Subcategory	N
Beauty & Personal Care	193	Humidifiers	17
Health & Household	159	Teeth Whitening Products	15
Home & Kitchen	148	Power Dental Flossers	14
Tools & Home Improvement	120	Sleep Sound Machines	12
Kitchen & Dining	112	Men's Rotary Shavers	11
Cell Phones & Accessories	81	Vacuum Sealers	11
Sports & Outdoors	77	Bug Zappers	10
Pet Supplies	62	Electric Back Massagers	10
Toys & Games	61	Cell Phone Replacement Batteries	9
Patio, Lawn & Garden	59	Light Hair Removal Devices	9
Electronics	57	Outdoor String Lights	9
Baby	42	Cell Phone Charging Stations	8
Office Products	30	Electric Foot Massagers	8

Fake Review Product Comparison

	Mean	SD	50%
Avg Rating			
Fake Review Products	4.40	0.51	4.50
All Products	4.23	0.59	4.30
Number of Reviews			
Fake Review Products	183.08	493.47	45.00
All Products	451.38	2619.02	59.00
Price			
Fake Review Products	33.36	44.96	23.99
All Products	44.69	154.80	20.99
Keyword Position			
Fake Review Products	21.41	16.11	16.00
All Products	28.18	17.32	23.00
Age (days)			
Fake Review Products	229.82	251.12	156.00
All Products	757.84	797.14	466.00
Sales Rank			
Fake Review Products	73292.27	151236.36	26200.50
All Products	89926.06	323028.92	21610.00

Set of comparison products are from same page of keyword search results.

What We Know About Fake Reviews on Amazon

From He, Hollenbeck & Proserpio (2022):

- Campaign starts: Immediate increase in ratings, reviews, and sales.
- Campaign stops: Immediate decrease in ratings, reviews, and sales.

From He, Hollenbeck, Overgoor, Proserpio, & Tosyali (2022):

- Can use review network structure to predict fake review purchasers
  - Accuracy = .858, AUC = .932
- We extend this by using review networks to label specific reviews.
  - 54% (median of 58%) for fake review purchasers

Outline of Talk

1. Simple model of welfare
2. Data and setting
3. **Empirical model of beliefs from ratings and demand**
4. Counterfactual Results

What Do We Need?

What determines the effect:

1. **Misinformation**: Shifts in demand with fake reviews.
   - Model mapping reviews to beliefs about quality.
   - Perceived-quality elasticity of demand.
2. **Mistrust**: Change in demand's response to ratings.
   - Model for how beliefs change with and without fake reviews.
3. **Pricing**: Price responses
   - Price elasticity of demand
   - Supply side

Model of Consumer Beliefs

Objective: Characterize how a *Bayesian* consumer forms expectations about quality from ratings.

1. How ratings are determined based on quality and fake reviews
   - Model assumptions
2. Beliefs about the prevalence of fake reviews
   - Surveys or rational expectations
3. Priors about the distribution of quality
   - Rational expectations, with estimates based on (1).

Model of Consumer Beliefs

A product of quality $q \in [0, 1]$ receives a positive review with probability:

$$p_{Fq} := \begin{cases} q & \text{if not a fake review purchaser (i.e., } F\text{=NFRP)} \\ \theta + (1 - \theta)q & \text{if a fake review purchaser (i.e., } F\text{=FRP)} \end{cases} .$$

Therefore the number of positive (N^+) and negative (N^-) reviews are distributed binomial:

$$P(N^+, N^- | q, F) = \binom{N^+ + N^-}{N^+} {p_{Fq}}^{N^+} (1 - p_{Fq})^{N^-} .$$

Model of Consumer Beliefs

A product of quality $q \in [0,1]$ receives a positive review with probability:

$$p_F q := \begin{cases} q & \text{if not a fake review purchaser (i.e., } F\text{=NFRP)} \\ \theta + (1-\theta) q & \text{if a fake review purchaser (i.e., } F\text{=FRP)}\ . \end{cases}$$

Therefore the number of positive (N^+) and negative (N^-) reviews are distributed binomial:

$$P(N^+, N^- | q, F) = \binom{N^+ + N^-}{N^+} {p_{Fq}}^{N^+} (1 - p_{Fq})^{N^-}\ .$$

Model of Consumer Beliefs

A product of quality $q \in [0,1]$ receives a positive review with probability:

$$p_{Fq} := \begin{cases} q & \text{if not a fake review purchaser (i.e., } F\text{=NFRP)} \\ \theta + (1-\theta)q & \text{if a fake review purchaser (i.e., } F\text{=FRP)} . \end{cases}$$

Therefore the number of positive (N^+) and negative (N^-) reviews are distributed binomial:

$$P(N^+, N^- | q, F) = \binom{N^+ + N^-}{N^+} {p_{Fq}}^{N^+} (1 - p_{Fq})^{N^-}.$$

Model of Consumer Beliefs

A product of quality $q \in [0,1]$ receives a positive review with probability:

$$p_{Fq} := \begin{cases} q & \text{if not a fake review purchaser (i.e., } F\text{=NFRP)} \\ \theta + (1-\theta)q & \text{if a fake review purchaser (i.e., } F\text{=FRP)} \end{cases}.$$

Therefore the number of positive (N^+) and negative (N^-) reviews are distributed binomial:

$$P(N^+, N^- | q, F) = \binom{N^+ + N^-}{N^+} p_{Fq}^{N^+} (1 - p_{Fq})^{N^-}.$$

Model of Consumer Beliefs

A product of quality $q \in [0,1]$ receives a positive review with probability:

$$p_{Fq} := \begin{cases} q & \text{if not a fake review purchaser (i.e., } F\text{=NFRP)} \\ \theta + (1-\theta)q & \text{if a fake review purchaser (i.e., } F\text{=FRP)}\ . \end{cases}$$

Therefore the number of positive (N^+) and negative (N^-) reviews are distributed binomial:

$$P(N^+, N^- | q, F) = \binom{N^+ + N^-}{N^+} {p_{Fq}}^{N^+} (1 - p_{Fq})^{N^-}.$$

Model of Consumer Beliefs

Posterior about quality given N^+ and N^- positive and negative reviews:

$$\begin{aligned} P(q|N^+,N^-) &= \sum_F P(F|N^+,N^-)P(q|N^+,N^-,F) \\ &= \sum_F P(F|N^+,N^-)\, P(N^+,N^-|q,F)\, P(q|F)\, /P(N^+,N^-|F) \end{aligned}$$

- $P(N^+,N^-|q,F)$: Binomial
- $P(F|N^+,N^-)$: Empirical or survey-based
- $P(q|F)$: Estimate via MLE. (Beta-distributed or non-parametric.)

Estimated P(q | F)

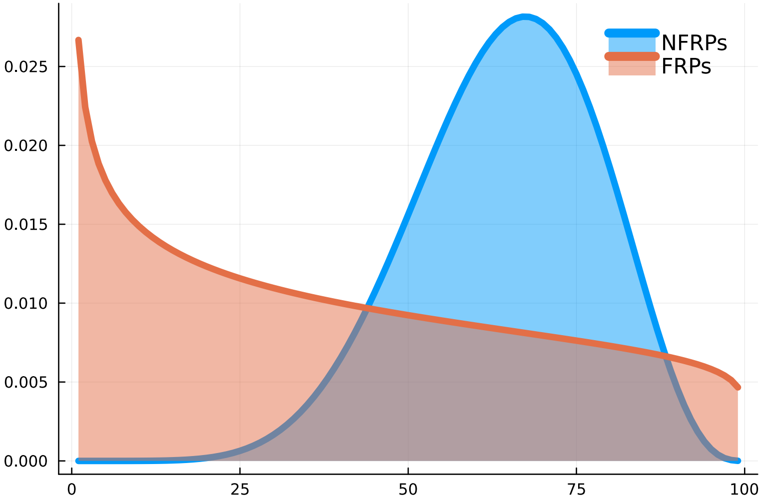

Non-parametric priors

Surveying Beliefs about Fake Reviews

Objective: Incentivized measure of beliefs about fake review prevalence.

1. Fraction of products: $P(F)$
2. Fraction given rating and number of reviews: $P(F|N^+, N^-)$
3. Fraction of fake reviews for fake review purchasers: θ

Number of quality responses from Prolific: 401.

Primary survey task

1. Each respondent selects 5 Amazon categories they shop in

2. Respondents are shown 10 product pages from these categories

3. Elicit perceived probability the product purchased fake reviews.
   - Incentive compatible with clear payoffs

Product Page

Please look at this product. Using the slider below, please select the percentage probability on a scale of 0 to 100 that the product purchases or has purchased fake reviews.

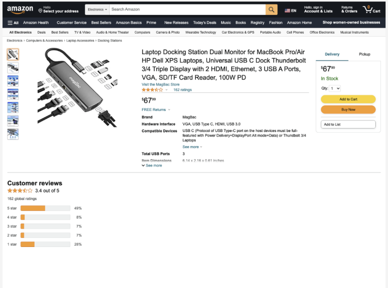

0% 20% 40% 60% 80% 100%

(You can only proceed after you make a selection on the slider.)

Rating and number of reviews are randomized (HTML)

Submit Probability with Clear Payoffs

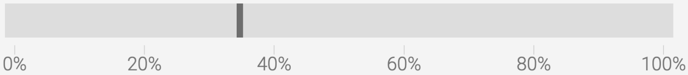

Your response indicates that you believe that there is a **35%** chance that the product purchased fake reviews.

If this product **did** purchase fake reviews, you will receive: **$0.35**.

If this product **did not** purchase fake reviews, you will receive: **$0.65**.

Distribution of Predictions on Fake Review Purchasing

Average prediction of $P(F)$ surprisingly close to rational expectations:

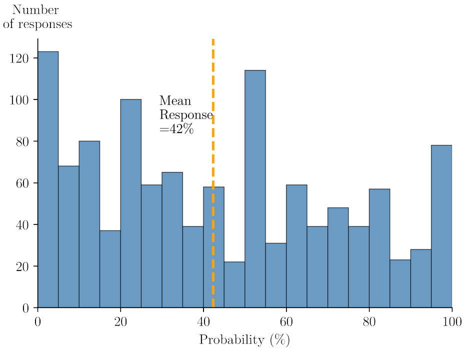

Fake Review Purchaser

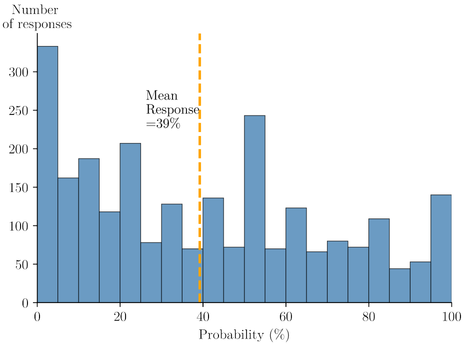

Non-Purchaser

Sanity check

$P(F|N^+, N^-)$: Beliefs Vary with Reviews

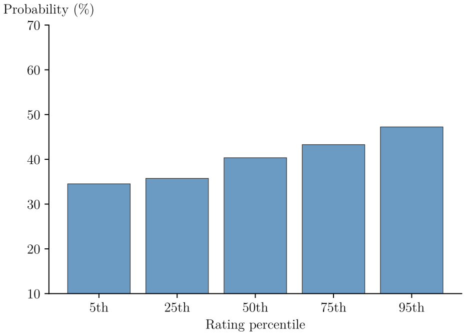

By rating

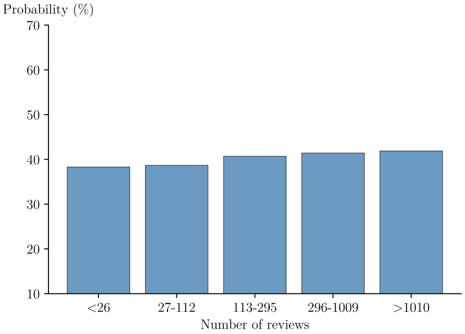

By number of reviews

$P(F|N^+,N^-)$: Beliefs Vary with Reviews

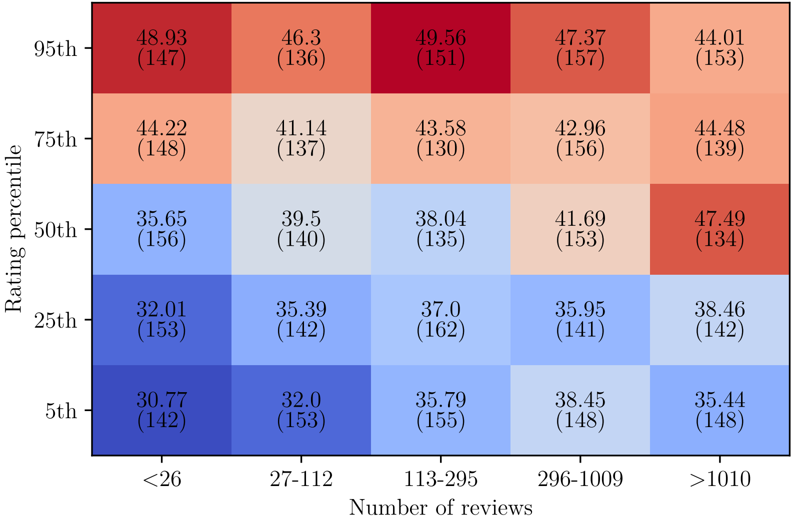

Extreme unimodal | Extreme bimodal | Heatmap in data

θ: Percent of Reviews that are Fake

Respondents underestimate the fraction of reviews that are fake (θ):

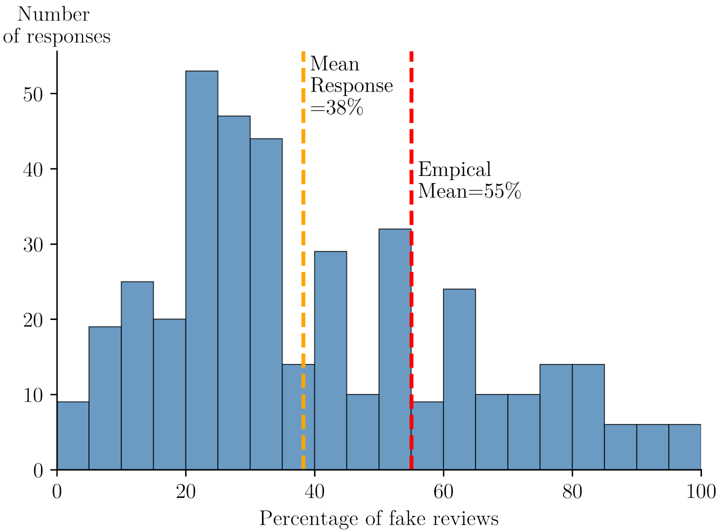

Demand Model

Utility for consumer i for product j in time t:

$$u_{ijt} = -\alpha \underbrace{p_{jt}}_{\text{Price}} + \beta_i \underbrace{E(q_{jt}|N^+_{jt}, N^-_{jt})}_{\text{Perceived Quality}} + \underbrace{X_{ijt}}_{\text{Other}} \psi + \underbrace{\mu_t + \lambda_j}_{\text{Week \& Product FEs}} + \underbrace{\xi_{jt}}_{\text{Demand shocks}} + \epsilon_{ijt}$$

- $E(q_{jt}|N^+_{jt}, N^-_{jt})$: Beliefs model transforms ratings into perceived expected quality.
    - Survey: $P(F|N^+, N^-)$, θ. Estimation: $P(q|F)$. Model: $P(N^+, N^-|q, F)$
- β_i: Allow heterogeneity in preference over quality.
- λ_j: Product FEs capture time-invariant product quality.
- μ_t: Time FEs captures seasonality in demand.
- X_{ijt}: Time-varying characteristics, including age and listing rank.

More Details Measuring Quantities

Demand Model - Results

Price	-0.051
	(0.034)
$E(q\|N^+, N^-)$	0.9
	(2.6)
σ	0.97
	(1.7)
Age	-0.042
	(0.069)
Listing Rank	-0.03
	(0.011)
ρ	0.14
	(0.49)
Product FEs	Yes
Week FEs	Yes
Gandhi-Houde IVs	Yes
Observations	37,501

Demand Model - Results

Price Elasticities	
Median Own	-0.96
Mean Own	-1.3
Mean Cross	0.46
Quality Elasticities	
Median Own	2.3
Mean Own	2.4
Mean Cross	-1.1

Outline of Talk

1. Simple model of welfare
2. Data and setting
3. Empirical model of beliefs from ratings and demand
4. **Counterfactual Results**

Counterfactuals

Objective: Assess the effect of fake reviews on sellers and consumers.

- Outcomes: quantities, prices, revenues, profits, consumer welfare

1. Effect of fake reviews on ratings and beliefs
2. Full equilibrium effect of fake reviews
3. Decomposing the channels

Product Positioning Dynamics

Deleting Fake Reviews Reduces Ratings

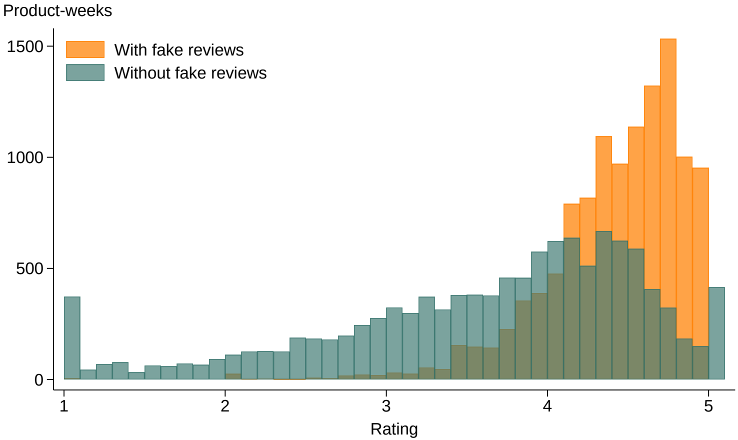

Expected Quality with Fake Reviews

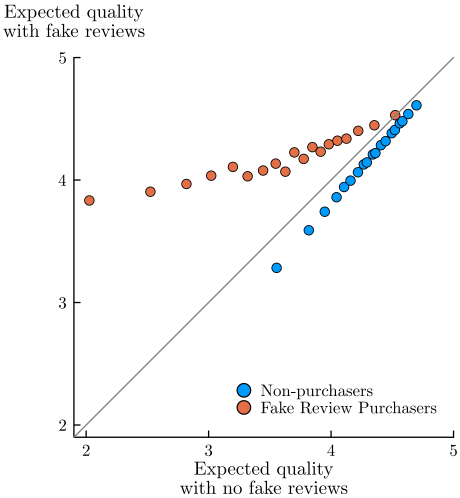

Changes in Quantities

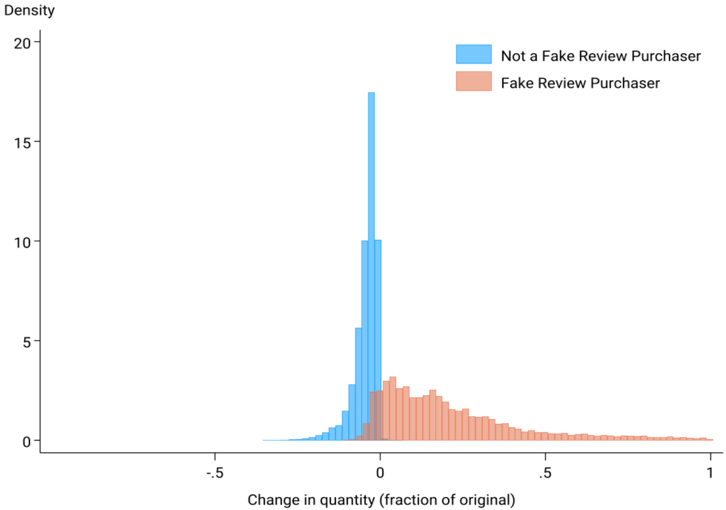

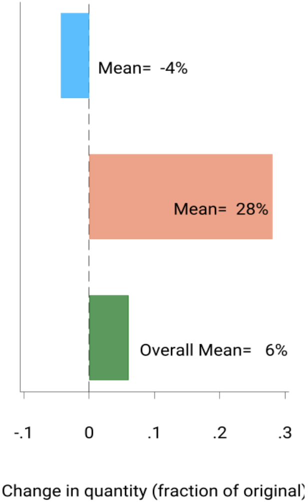

Changes in Prices

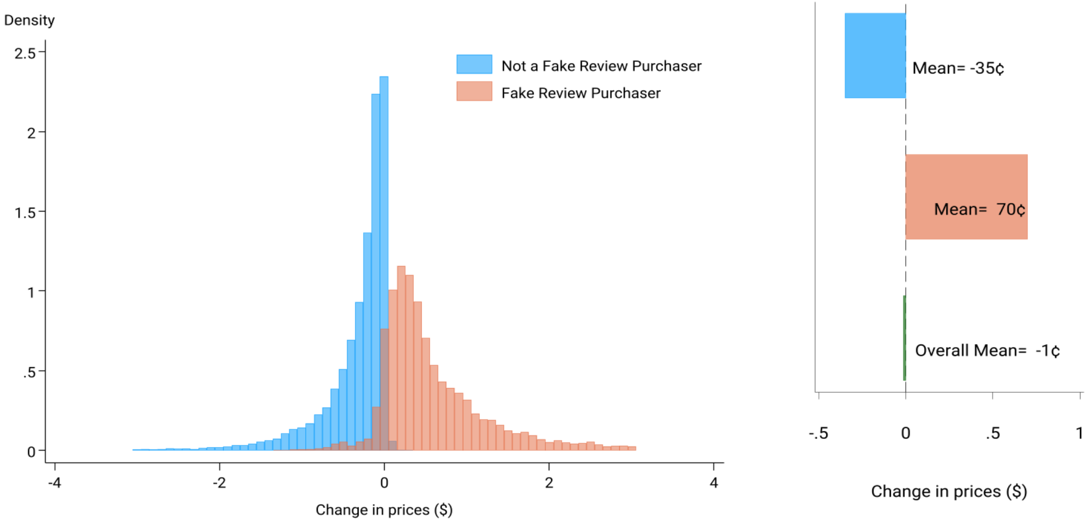

Changes in Profits

Additional Profit per Fake Review

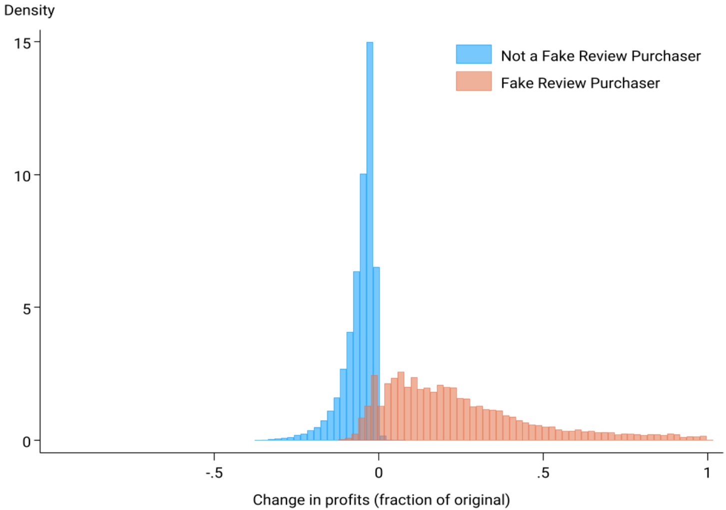

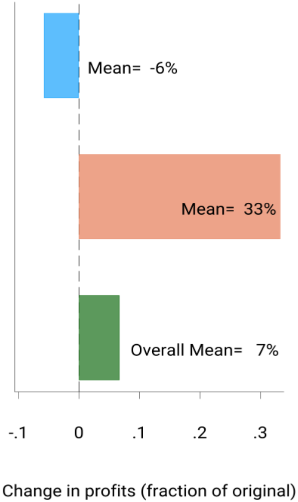

Changes in Welfare (*Preliminary)

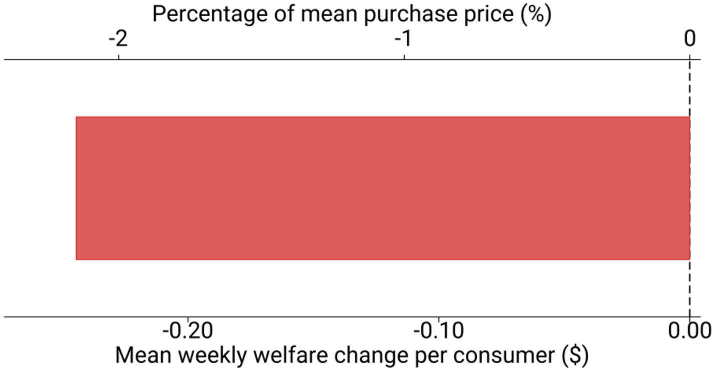

Counterfactuals

Objective: Assess the effect of fake reviews on sellers and consumers.

- Outcomes: quantities, prices, revenues, profits, consumer welfare

1. Effect of fake reviews on ratings and beliefs
2. Full equilibrium effect of fake reviews
3. Decomposing the channels

Product Positioning Dynamics

Counterfactuals - Isolating the Mechanisms

	No FR	Misinfo	Mistrust	Misinfo+Mistrust	
				Fixed prices	Floating prices
Welfare ($)	38,491,022	38,202,942	38,654,694	38,179,622	38,404,980
Platform revenue ($)	4,026,278	4,119,159	3,912,182	4,002,119	4,017,978
FRP average prices ($)	26.78	27.56	26.62		27.48
NFRP average prices ($)	30.95	30.70	30.87		30.60
FRP sales (units)	318,071	376,876	303,841	388,436	369,208
NFRP sales (units)	1,073,026	1,052,593	1,054,482	1,010,719	1,032,205
FRP profits ($)	6,732,360	8,065,500	6,402,327	7,936,544	7,878,878
NFRP profits ($)	24,975,683	24,340,059	24,492,908	23,721,193	23,808,138

Misinformation alone harms consumers, but mistrust alone can benefit them by increasing price competition (similar to lit on information disclosure (Saeedi & Hopenhayn (2022), Vatter (2021))

Counterfactuals - Isolating the Mechanisms

	No FR	Misinfo	Mistrust	Misinfo+Mistrust	
				Fixed prices	Floating prices
Welfare ($)	38,491,022	38,202,942	38,654,694	38,179,622	38,404,980
Platform revenue ($)	4,026,278	4,119,159	3,912,182	4,002,119	4,017,978
FRP average prices ($)	26.78	27.56	26.62		27.48
NFRP average prices ($)	30.95	30.70	30.87		30.60
FRP sales (units)	318,071	376,876	303,841	388,436	369,208
NFRP sales (units)	1,073,026	1,052,593	1,054,482	1,010,719	1,032,205
FRP profits ($)	6,732,360	8,065,500	6,402,327	7,936,544	7,878,878
NFRP profits ($)	24,975,683	24,340,059	24,492,908	23,721,193	23,808,138

Platform profits from misinformation, but the long-run cost of allowing fake reviews manifests as lower trust

Discussion and Conclusion

- Fake reviews are widespread and is of growing interest to regulators
  - First empirical examination that considers equilibrium effects, including through price competition and trust.
- Fake reviews are responsible for large changes in quantities, prices, revenues, profits, and welfare.
  - Harm consumers
  - Substantially benefit purchasers and harm honest products
  - Increase profits for the platform
  - Equilibrium changes in pricing and trust are important
- Much more to be done: endogenous purchasing, dynamics, heterogeneous sophistication, interaction between platforms.

Thanks!

- Email: ashvin.gandhi@anderson.ucla.edu
- Email: brett.hollenbeck@anderson.ucla.edu

Short-Run Effects of Fake Review Campaigns

From He, Hollenbeck & Proserpio (2022)

- In addition to directly observing who buys fake reviews, a unique aspect of this data is the panel on firm outcomes
- Sharp and immediate increase in avg. rating, weekly # reviews, and sales

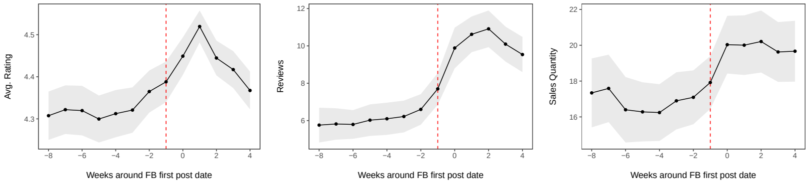

Model Back

Background: Long Term Outcomes

From He, Hollenbeck & Proserpio (2022)

- Track outcomes after rating manipulation stops
- Sharp and immediate decrease in sales rank, avg. rating, and weekly # reviews

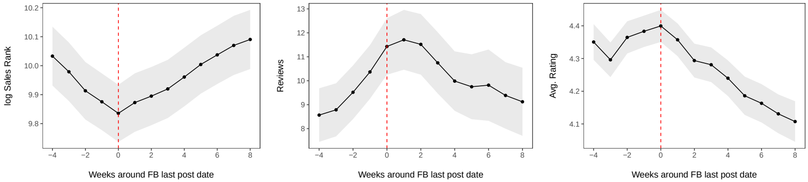

Back

Model - Beliefs - Estimating θ^F

To estimate θ_j^F, we rely on He, Hollenbeck, Overgoor, Proserpio, & Tosyali (PNAS 2022), who develop a model to predict what products buy fake reviews with high accuracy (Accuracy = .858, AUC = .932). We build out the product-reviewer network and:

1. Classify all products as fake review products (FRPs) or not (NFRP)
2. Classify all reviewers based on leaving 5-star reviews for multiple FRPs
3. Classify all reviews as fake if 5-star and left by fake reviewer
4. This provides an *estimate* of θ_j^F

For FRPs, the average share of fake reviews is 54% (median of 58%.)

▸ Back

Simple Economics of Fake Reviews

How much does a seller pay the reviewer for 1 fake review?

$$P(1 + \tau + F_{PP}) + Commission$$

Where:

- P = list price
- τ=sales tax
- F_{PP} = PayPal fee
- Commission is generally zero but sometimes $5-10

Simple Economics of Fake Reviews

How much do they get from Amazon for the fake sale?

$$P(1-c)$$

Where c is Amazon's commission on each sale. So the difference in payments or net cost of 1 review is:

$$P(1+\tau+F_{PP}) - P(1-c) = P(\tau + F_{PP} + c)$$

And with the production cost of the product (MC), the full cost of 1 fake review is:

$$Cost = MC + P(\tau + F_{PP} + c)$$

Simple Economics of Fake Reviews

Define the seller's markup λ such that $P(1-\lambda) = MC$ (i.e. $\lambda = \frac{P-MC}{P}$)

$$Cost = P(1 - \lambda + \tau + F_{PP} + c)$$

The benefit of 1 fake review is a function of how many organic sales it creates Q_o, markup, and commission:

$$Benefit = Q_o P(\lambda - c)$$

Define Q_o^{BE} as breakeven quantity: the # of sales necessary to exactly justify buying 1 fake review:

$$Q_o^{BE} = \frac{1 - \lambda + \tau + F_{PP} + c}{\lambda - c}$$

Simple Economics of Fake Reviews

- Sales tax $\tau = .0656$[1]
- Paypal fee $F_{PP} = 2.9\%$
- Amazon commission c is either 8% or 15% in almost all cases.[2]

The result is:

$$Q_o^{BE} = \frac{1.175 - \lambda}{\lambda - .08}$$

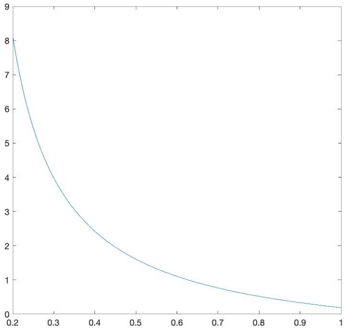

[1] https://taxfoundation.org/2020-sales-taxes/,simple average over states
[2] https://sellercentral.amazon.com/gp/help/external/200336920

Simple Economics of Fake Reviews

Implications:

1. Economics of rating manipulation potentially quite favorable for sellers
2. Lower quality products need far fewer sales to justify a fake review
   - Imagine two products that both list a price of $25. Product A costs $15 to produce and product B costs $20 to produce because A is of lower quality than B.
   - For product A: $Q_0^{BE} = 2.4$
   - For product B: $Q_0^{BE} = 8.1$
3. Unlikely we would see fake negative reviews for competitors
   - Marginal cost much higher
   - Marginal benefit presumably much lower

Back

Extreme Bimodal Rating Distributions

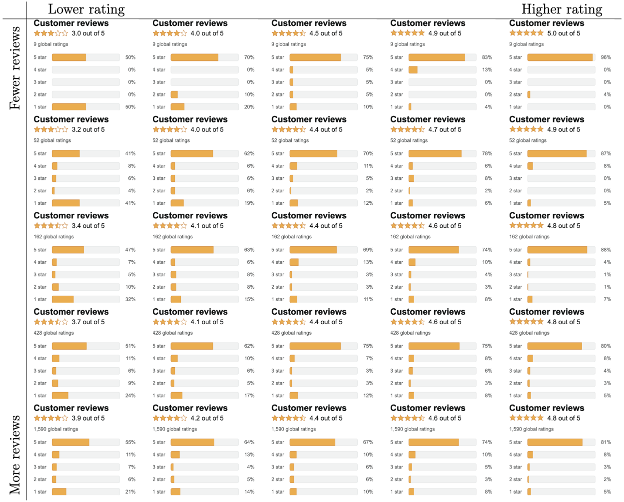

Extreme Unimodal Rating Distributions

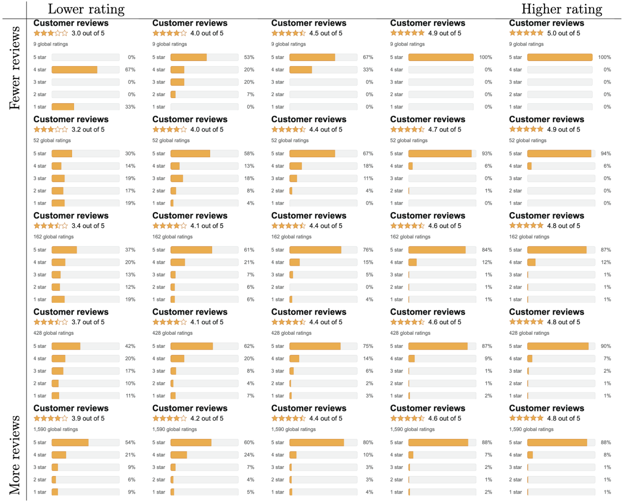

Observed $P(F|N^{+},N^{-})$

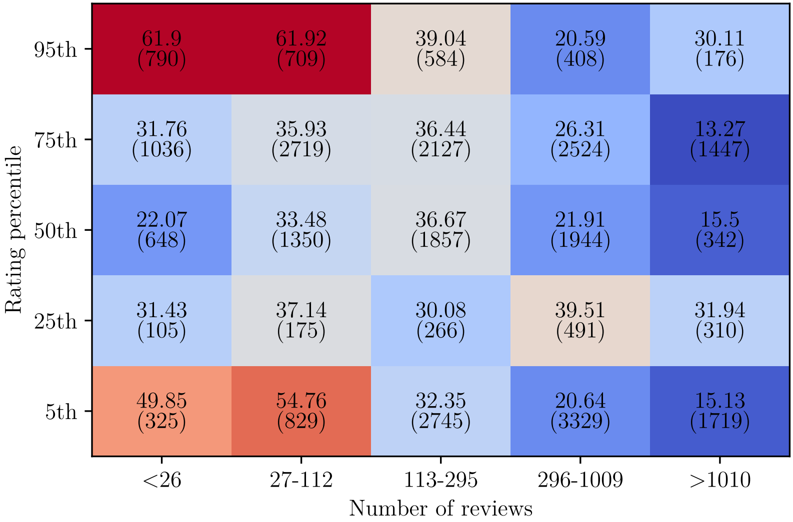

Sanity Check: Amazon Gift Card

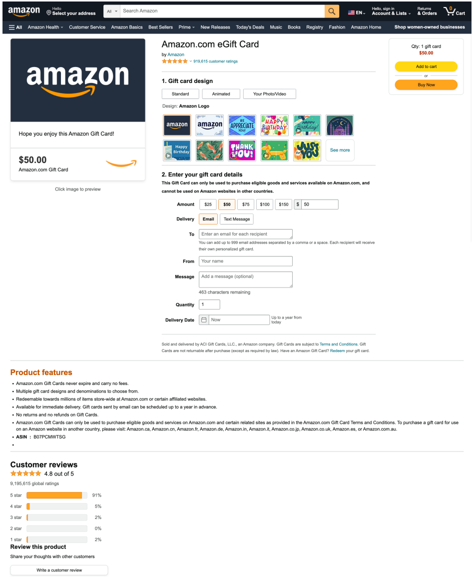

Sanity Check: Amazon Gift Card

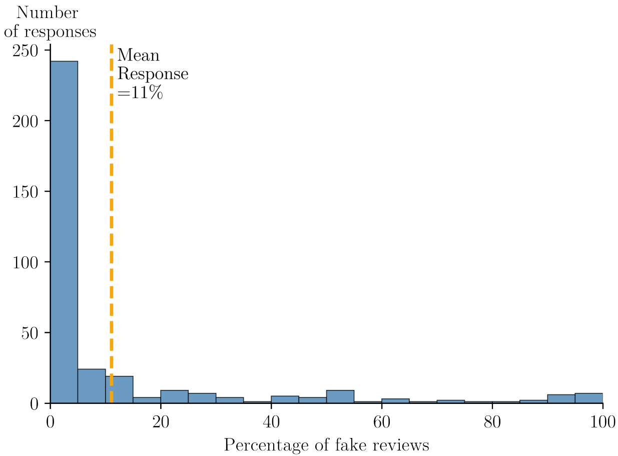

50% of the respondents select 0.

Non-Parametric Priors

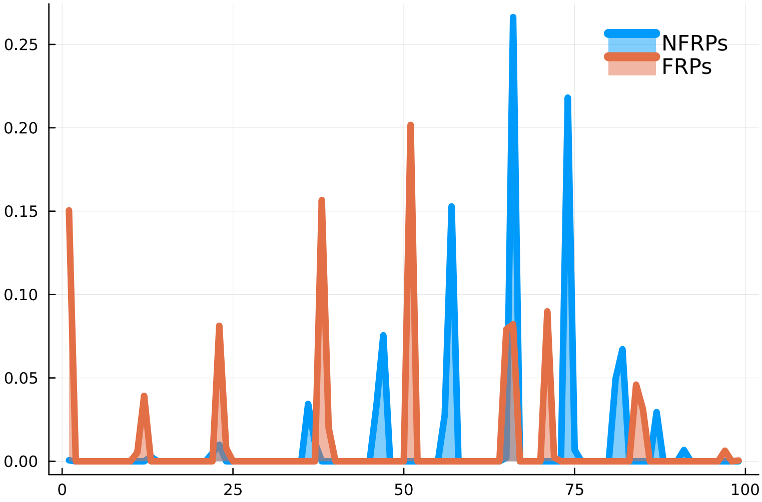

Back

Demand Model - Implementation

$$u_{ijt} = -\alpha p_{jt} + \beta E(q_{jt}|N_{jt}^{+}, N_{jt}^{-}) + X_{ijt}\psi + \mu_t + \lambda_j + \xi_{jt} + \epsilon_{ijt}$$

Other components of demand model

- Markets determined by up to 10 products that frequently co-occur in keyword search results.
- Nested logit structure on outside good

Endogeneity:

- Price: use Gandhi-Houde IVs - constructed from competitor characteristics
- Fixed effects for product capture time-invariant product quality
- Fixed effects for week capture seasonality

Measuring Quantities

- We observe sales rank daily for all products
- Calculate quantities following He & Hollenbeck (2021):
  1. Observe inventories for products with fewer than 1000 units available.
     - Most products in most weeks.
  2. Collect inventory data every 2 days during sample period.
  3. Compute daily sales using observed drops in inventory.
  4. Estimate relationship between daily sales and sales rank to interpolate (sometimes extrapolate) missing data.

Counterfactuals - Expected Quality

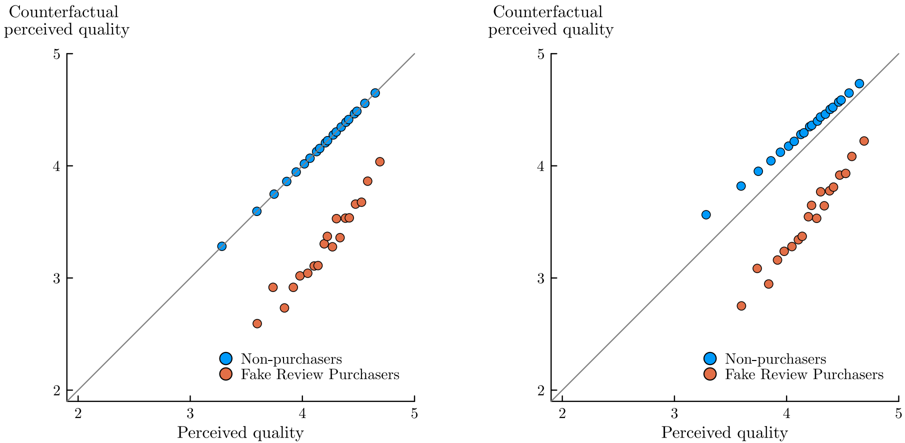

Left: Discreet deletion of FRs. Right: Consumers also update beliefs.

Hedonic Model of Product Position

L1.Log Shares	0.494***	0.531***	0.420***	0.515***
	(13.30)	(13.04)	(12.57)	(12.54)
L2.Log Shares	0.274***	0.328***	0.258***	0.290***
	(7.76)	(8.10)	(7.50)	(7.10)
L1.Log N. Good Reviews	0.128***			
	(9.38)			
L2.Log N. Good Reviews	0.0820***			
	(5.95)			
L1.Cumulative rating		0.123***		0.166***
		(3.77)		(4.59)
L2.Cumulative rating		0.126***		0.140***
		(3.95)		(3.98)
L1.Weekly rating			0.0471***	-0.0222
			(4.09)	(-1.54)
L2.Weekly rating			0.0298**	-0.0284*
			(2.58)	(-2.01)
L1.Log Cumulative N. Reviews		0.0492		0.00762
		(1.95)		(0.29)
L2.Log Cumulative N. Reviews		0.0101		-0.0169
		(0.42)		(-0.68)
L1.Log Weekly N. Reviews			0.0901***	0.0604***
			(5.50)	(3.34)
L2.Log Weekly N. Reviews			0.0918***	0.0863***
			(5.76)	(4.95)
Sponsored	0.563***	0.552***	0.574***	0.555***
	(8.45)	(8.30)	(8.63)	(8.35)
Log Age	0.260	0.407**	0.503***	0.533***
	(1.74)	(2.71)	(3.35)	(3.56)
Constant	-1.170**	-1.743***	-1.551***	-2.082***
	(-2.86)	(-4.26)	(-3.81)	(-5.09)
Observations	73933	73933	73933	73933

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Additional Profits from Fake Review Purchase

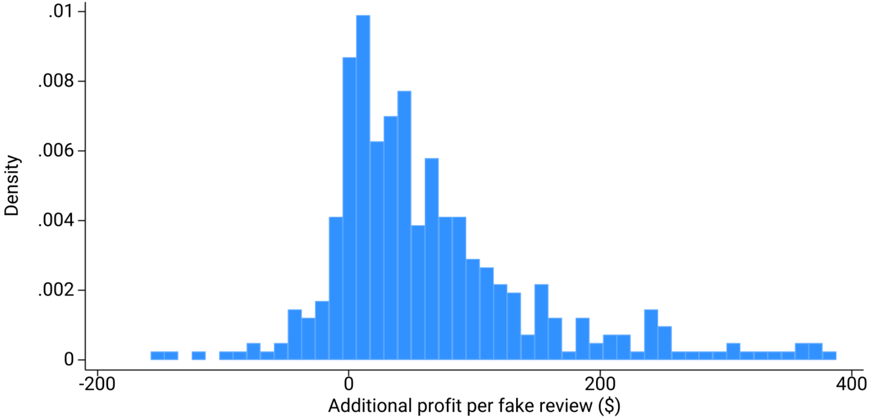

- Mean=$64, Median = $43

Heterogeneity in Market Share by Share of Fake Reviews

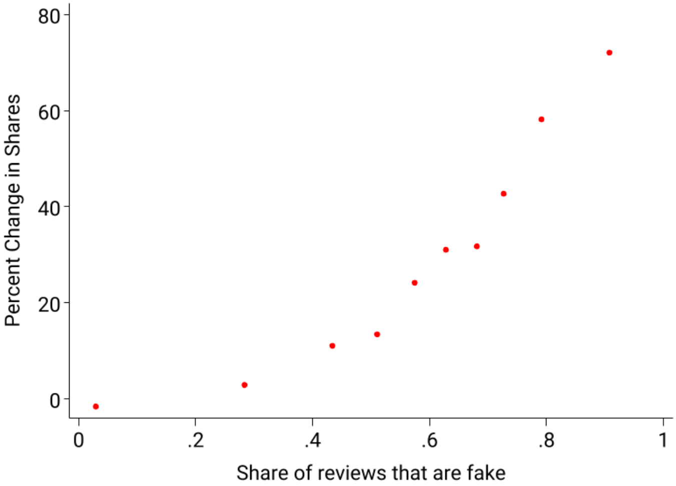

Heterogeneity in Prices by Share of Fake Reviews

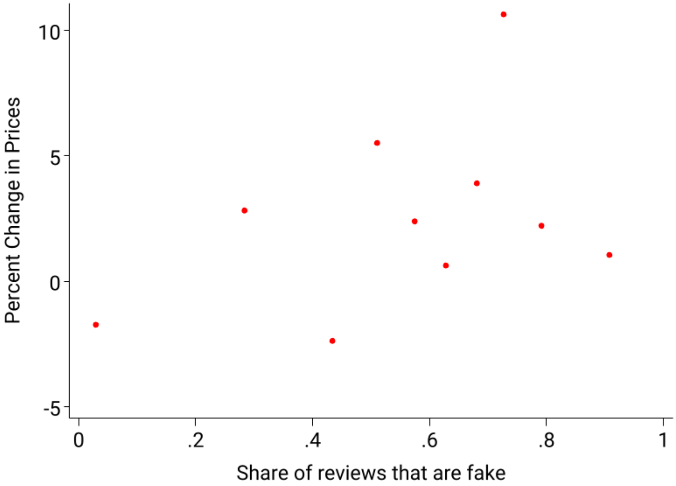

Heterogeneity in Profits by Fake Reviews' Effect on Expected Quality

Changes in Quantities

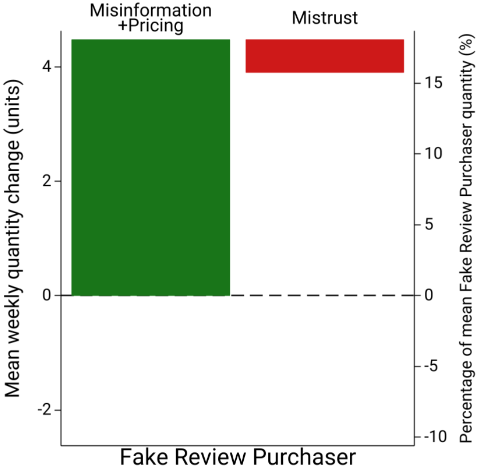

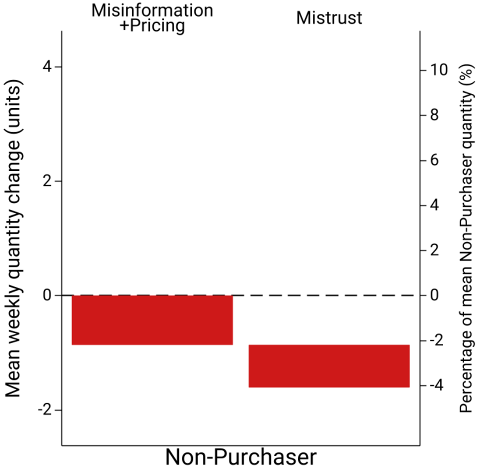

Changes in Revenue

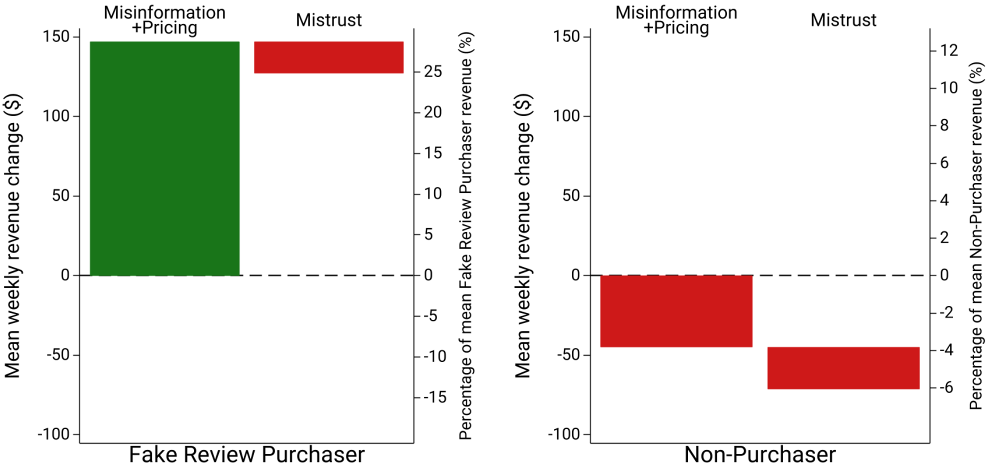

Changes in Profits

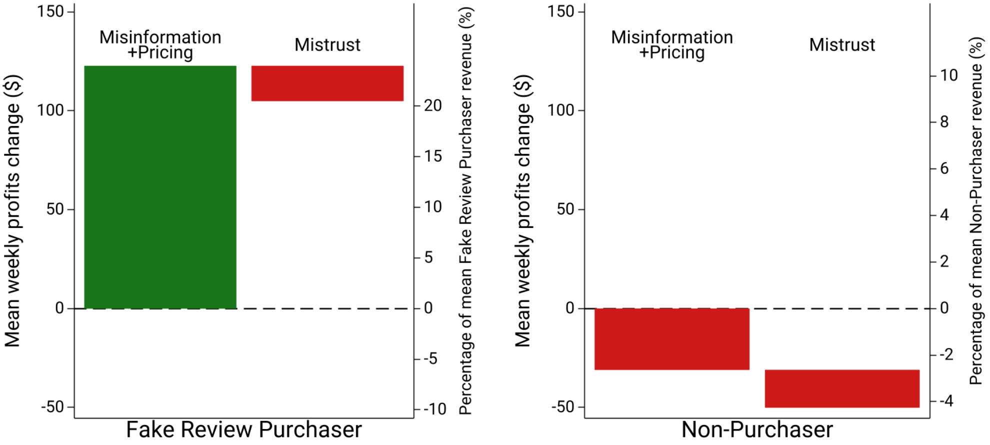

Counterfactuals - Computing Welfare ex post

For all counterfactuals we compute welfare at the actual estimated product quality, not the expected quality used in demand, which will differ for fake review products.

- We compute experience utility $\tilde{u}_{ijt}$ with an offset term that depends on the discrepancy between perceived and true qualities and the estimated coefficient on quality.

$$\Delta q := q_{perceived} - q_{true}$$
$$\tilde{u}_{ijt} = u_{ijt} - \beta_1 \Delta q_{ijt}$$

The welfare for consumer i in market t is then

$$\begin{aligned} W_{it} &= E_\epsilon[u_{ij^*t}] - E_\epsilon[\Delta q_{ij^*t}] \\ &= \bar{W}_{it} - \sum_{J_t} s_{ijt}(\beta_1 \Delta q_{ijt}), \end{aligned}$$

where j^* is chosen based on perceived quality, and $\bar{W}_{it}$ is the welfare evaluated using decision utility.

Back