Bonus or Not? Learn to Reward in Crowdsourcing

Ming Yin, Yiling Chen

John A. Paulson School of Engineering and Applied Sciences, Harvard University

Introduction

Monetary reward is a primary type of incentives on crowdsourcing platforms like Amazon Mechanical Turk, and recent studies showed that the placement of performance-contingent rewards (e.g. bonuses) ***can*** affect work quality! But in a working session…

How do workers react to bonuses provided in selected tasks?

✚

What is the trade-off between quality and costs?

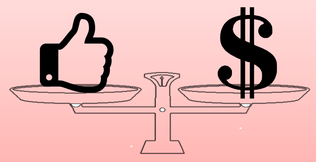

Whether and when to provide a bonus for a worker in a working session?

Our approach: Input-output Hidden Markov Model ✚ Requester Utility Function ➔ Online Decision Making

An Algorithmic Approach

1

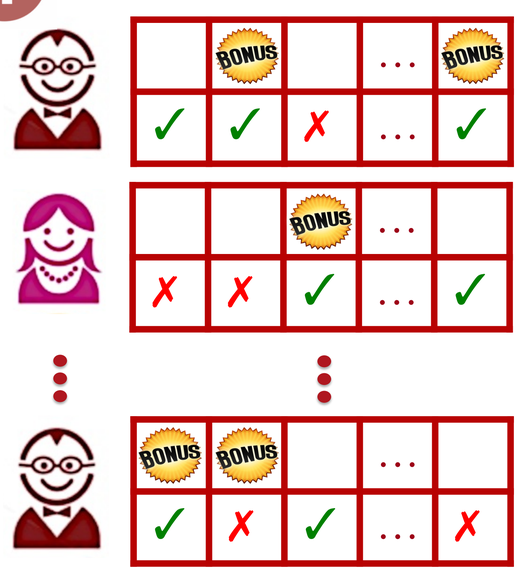

Training Data

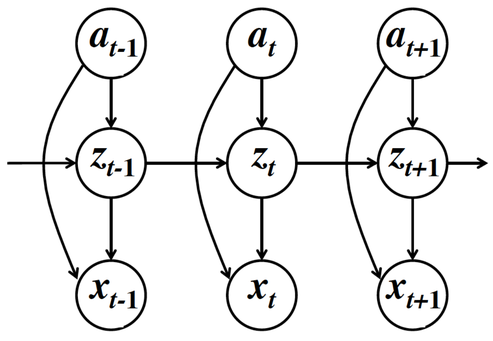

IOHMM

Use EM algorithm to learn a worker behavior IOHMM from the training data.

2

$$U = w_l N_{LQ} + w_h N_{HQ} - c N_{bonus}$$

Requester's Utility

3

Online Decision Making (equivalent to solve POMDP)

$$a_{t_c+1} = \underset{a \in \{0,1\}}{\operatorname{argmax}}\, EU_{max}(\boldsymbol{b}(t_c), a, T - t_c)$$

$EU_{\max}(\boldsymbol{b}, a, l)$ is the maximum expected utility a requester can obtain for the next l tasks, when the current state belief is $\boldsymbol{b}$, the input for the next task is a.

Heuristics

- n-step look-ahead
- MLS-MDP
- Q-MDP

MTurk Experiment & Simulation

MTurk experiment: Word puzzle game (Training: 50 workers, 20% random bonus; Testing: 6 treatments × 50 workers)

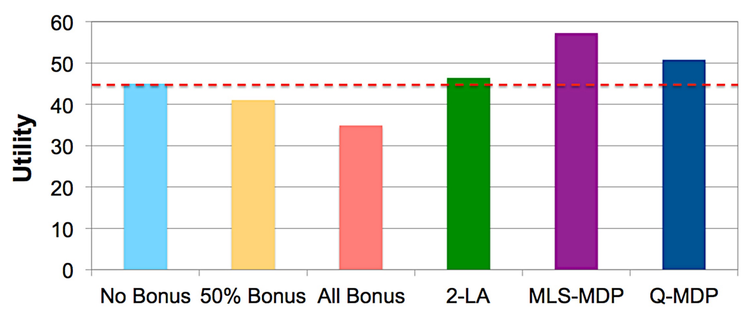

# HQ answers	300	331	349	322	395	368
Cost ($)	0	8.60	17.45	1.90	2.00	4.40

Worker	Inputs & Outputs in the Working Session								
A	Bonus?	✗	✗	✗	✗	✗	✗	✗	✗	✗
	High-quality?	1	1	1	1	1	1	1	1	1
B	Bonus?	✗	✓	✓	✓	✓	✓	✓	✓	✓
	High-quality?	0	0	0	0	0	0	0	0	0
C	Bonus?	✗	✓	✓	✓	✓	✓	✗	✗	✗
	High-quality?	0	0	0	1	1	1	1	1	1
D	Bonus?	✗	✗	✗	✗	✓	✓	✓	✗	✗
	High-quality?	1	1	0	0	1	1	1	1	0

O	V	H	L	E	G	T	Y	J	E	Z	J
P	V	L	M	Q	A	K	L	M	L	P	Q
M	A	A	Y	W	M	T	A	K	U	G	M
P	G	E	Q	G	E	G	M	X	U	A	Q
Y	A	K	B	C	W	H	R	S	S	M	K
Q	M	K	G	A	M	E	A	N	K	E	I
A	E	G	D	E	M	K	L	L	M	J	Y
E	M	G	M	A	N	B	E	A	X	F	P
M	X	A	G	Z	K	I	G	G	A	M	E
W	G	D	B	D	S	Y	P	Z	A	E	X
Z	P	W	Y	W	J	E	H	T	M	M	M
R	K	I	K	M	I	O	V	N	L	V	E

More high-quality work ✚ Lower cost ➔ Higher requester utility

- Differentiate "diligent" workers from "lazy" workers and strategically focus on incentivize "lazy" workers
- Timing: Keep incentivizing until stabilized good performance / Add extra incentives in time to bring back good effort.

Simulation: Two worker behavior models (Training: 3000 workers, 20% random bonus; Testing: 6 treatments × 100 workers)

Our approach is robust against different worker behavior & worker population in improving requester utility

Worker's accuracy changes from acc^l to acc^h when bonus is provided. Each worker draws acc^l and acc^h from:

- A set of discrete values
- Beta distributions

Model 1: Two Capability Levels

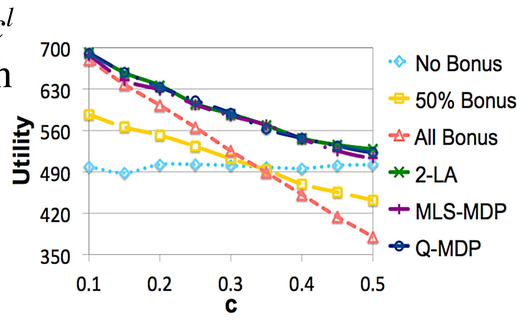

Workers compare the current payment with the reference and decide the work quality:

- α (skill); β (responsiveness to rewards)
- Mixing workers with different α & β

Model 2: Reference Payment Levels