Crowdsourcing Contests

Ruggiero Cavallo
Microsoft Research NYC

CS286r: November 5, 2012

What is crowdsourcing (for today)?

- Principal seeks production of a good; multiple agents produce; principal obtains value commensurate with *highest* quality good.

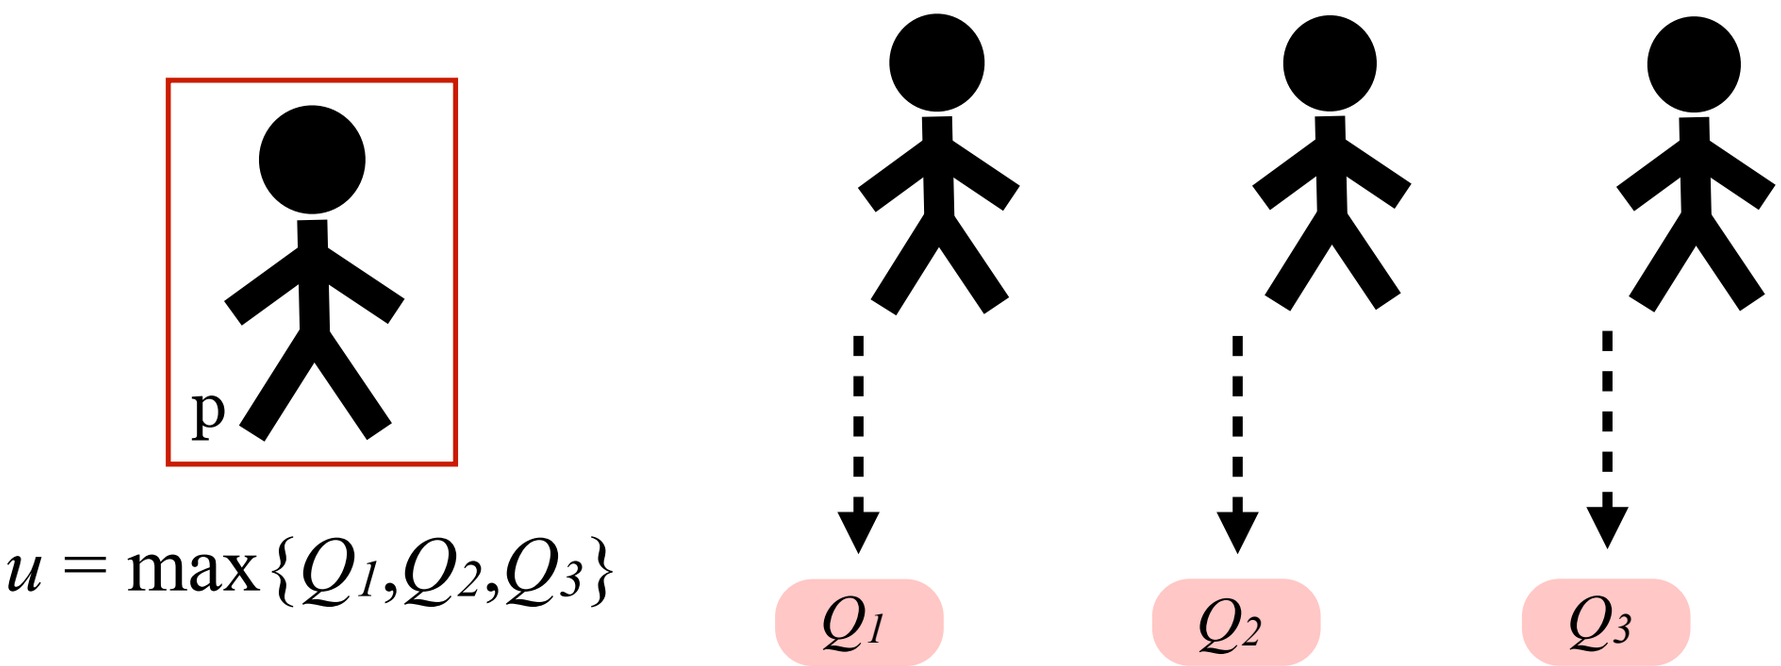

What is crowdsourcing (for today)?

- Principal seeks production of a good; multiple agents produce; principal obtains value commensurate with *highest* quality good.
- Examples: logo design, web page design, software development, question answering.
- Getting popular on the web – 99designs, Taskcn, Topcoder, Innocentive, CrowdCloud, CrowdFlower, ... Amazon Mechanical Turk, Yahoo! Answers
- $ And stakes are growing: 99designs.com has paid out over $40,000,000 to community of 180K designers

- number of producers can be very large
- traditionally: only one wins and obtains a "prize"

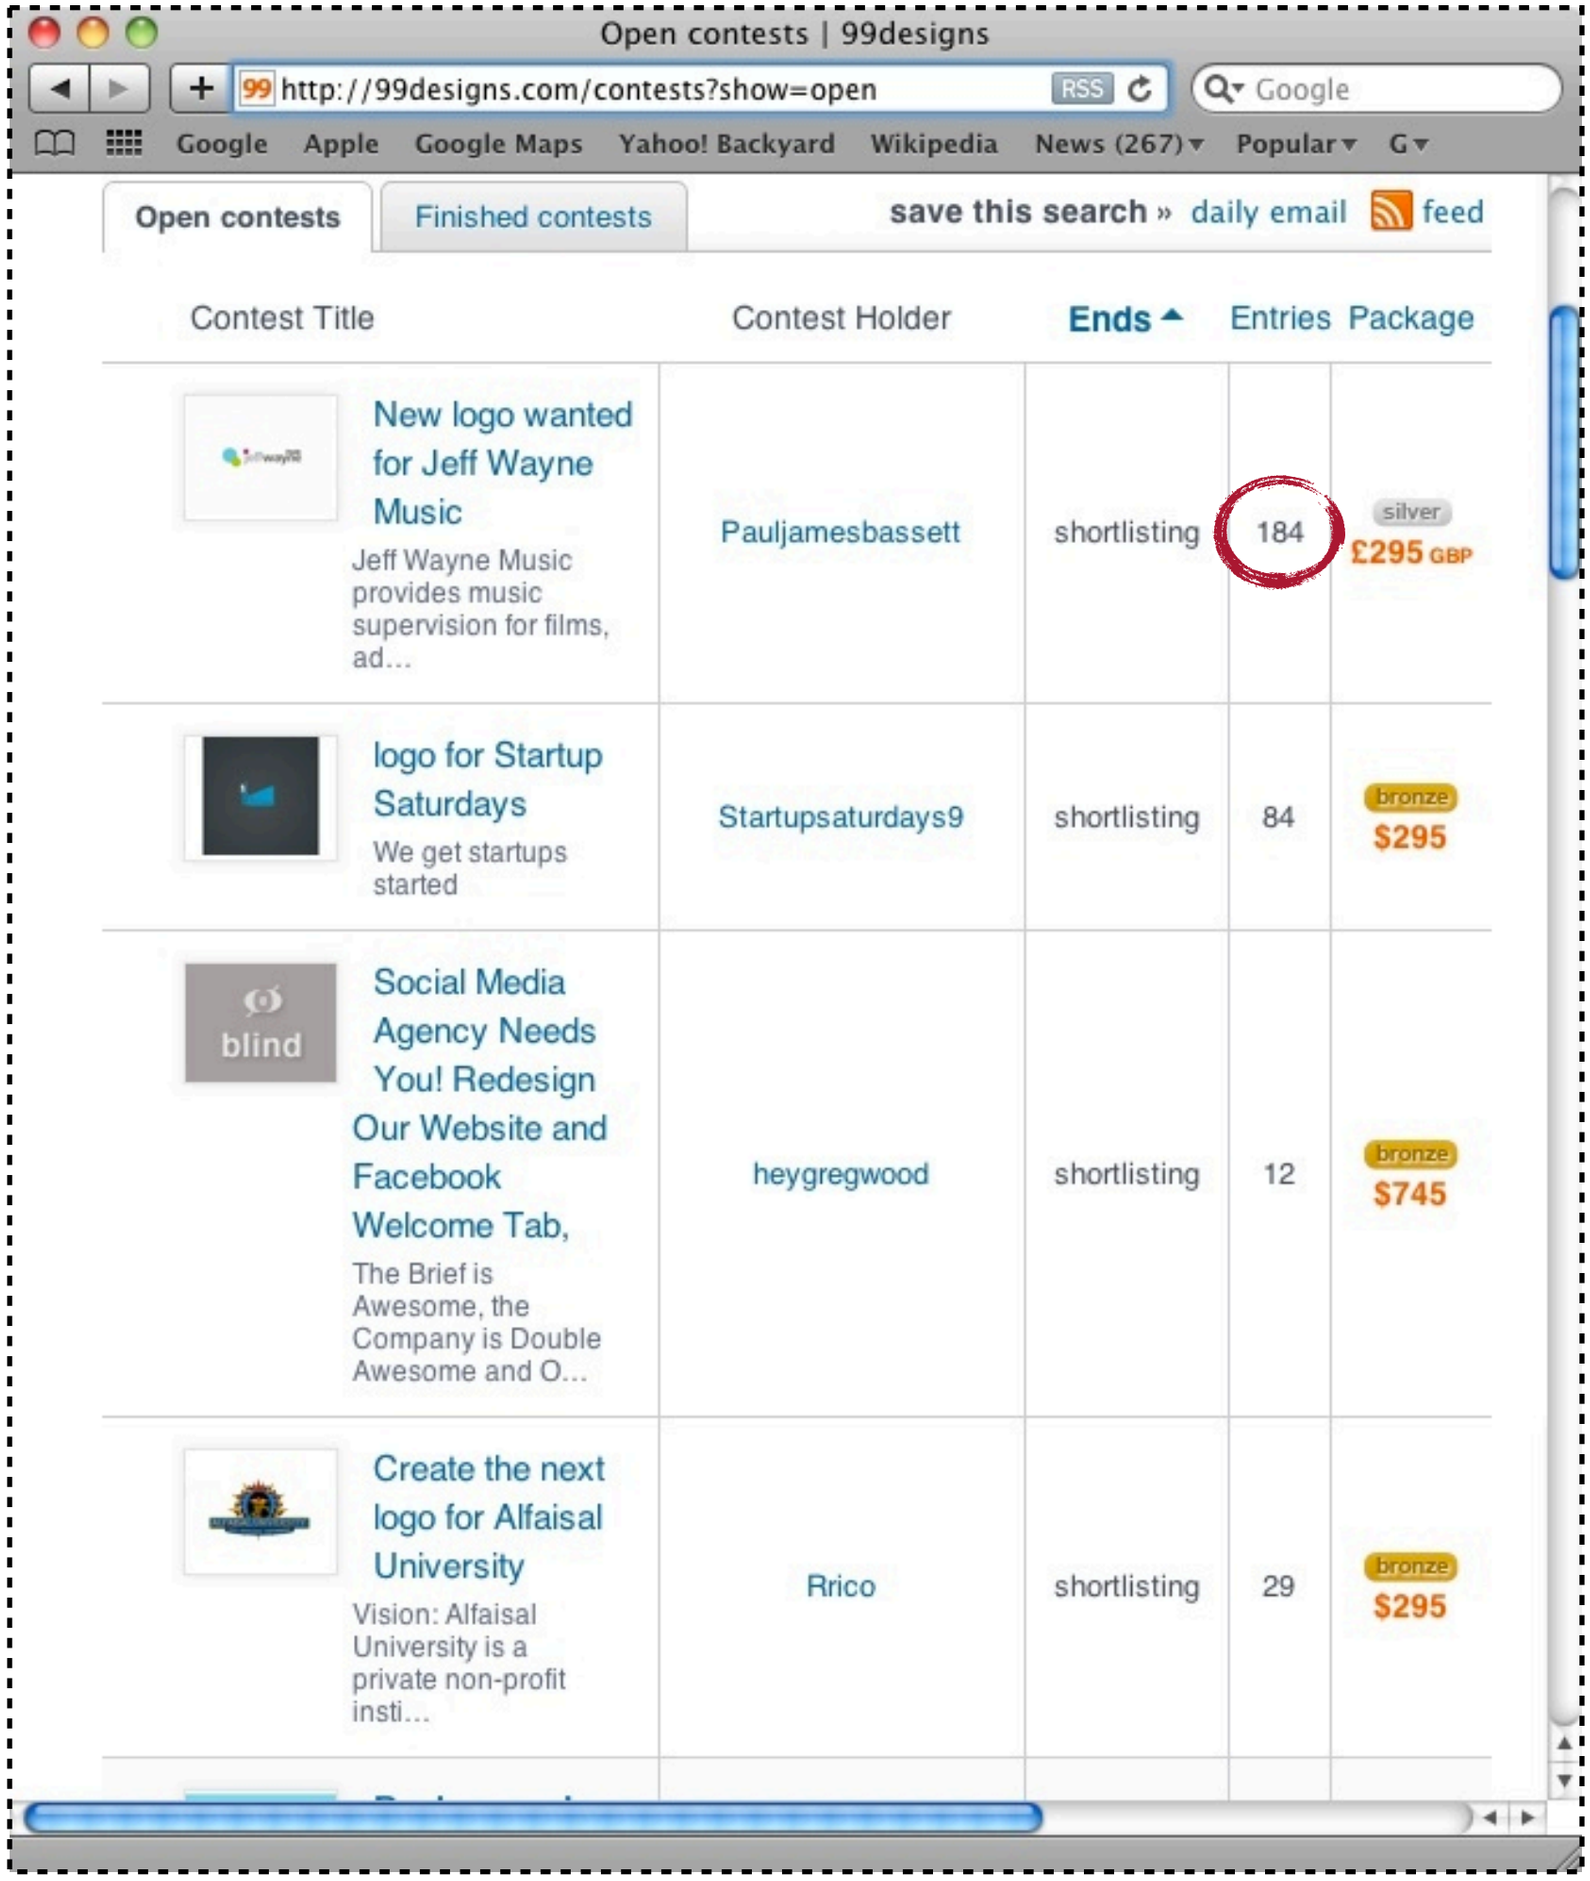

Backing up a little...

- This isn't quite new, or primarily internet-based.
  - Defense contracting (competitors build prototypes, competing for large contract).
  - X prize (spacecraft, fuel efficient car, tricorder).
  - American Idol?

Main existing theory

- Contest design in economics (just a sampling):
  - [Fullerton and McAfee, 1999]
  - [Moldovanu and Sela, 2001, 2006]
- More recently, specifically motivated by online crowdsourcing:
  - [DiPalantino and Vojnovic, 2009]
  - [Chawla, Hartline, and Sivan, 2012]
  - [Archak and Sundararajan, 2009]
  - [Cavallo *(me)* and Jain, 2012]

Auctioning Entry Into Tournaments [Fullerton and McAfee, 1999]

- Research tournaments, where participants bear *fixed cost* plus cost of research effort.
- Principal seeks to maximize best submission net of prize paid out.
  - Cost of obtaining a given equilibrium quality level is minimized with 2 participants.
  - To get the *best* participants, conduct a preliminary all-pay auction, which implicitly reveals highest-skilled agents.

Auctioning Entry Into Tournaments

[Fullerton a

99designs.com now similarly has "qualifying" and "final" rounds (where principal chooses up to 6 finalists).

- Research to
  bear *fixed cost* plus cost of research effort.
- Principal seeks to maximize best submission net of prize paid out.
  - Cost of obtaining a given equilibrium quality level is minimized with 2 participants.
  - To get the *best* participants, conduct a preliminary all-pay auction, which implicitly reveals highest-skilled agents.

Crowdsourcing and All-Pay Auctions [DiPalantino and Vojnovic, 2009]

- Agents (workers) have private skill, drawn from common-knowledge distribution, which determines how costly it is to produce at a given quality level.
- Agents choose among **multiple contests** to participate in, and choose effort level.
- In each contest, agent with highest quality submission receives a prize.
  - Model equilibrium participation rates as a function of prize-value, compare with empirical data from TaskCN.

Optimal Crowdsourcing Contests [Chawla, Hartline, and Sivan, 2012]

- Adopt model of [DiPalantino and Vojnovic, 2009] – analogous to all-pay auction, since all agents pay and only highest "bidder" (quality submitter) obtains the "good" (prize).
- **Principal-optimal** mechanism design, seeking to maximize either sum of qualities or max quality.
  - For *sum-of-qualities* goal: approximation result (3.164-approx).
  - For *max-quality* goal: winner-take-all is optimal "fixed-prize" format; more messy characterization for the general case.

- Almost all previous papers consider the *principal's perspective*: how to elicit optimal submission (or sum of submission qualities).
- All (i.e., both of) the main previous computer science papers consider *deterministic* production.

▶ Rest of the lecture: design of an *efficient* crowdsourcing mechanism with stochastic production [Cavallo and Jain, 2012].

- optimally trade off benefit to principal with costs to agents

When does crowdsourcing make sense?

- Two key factors:
  1. Uncertain quality of production
  2. Impatience / deadline
- Otherwise better to just order production sequentially.

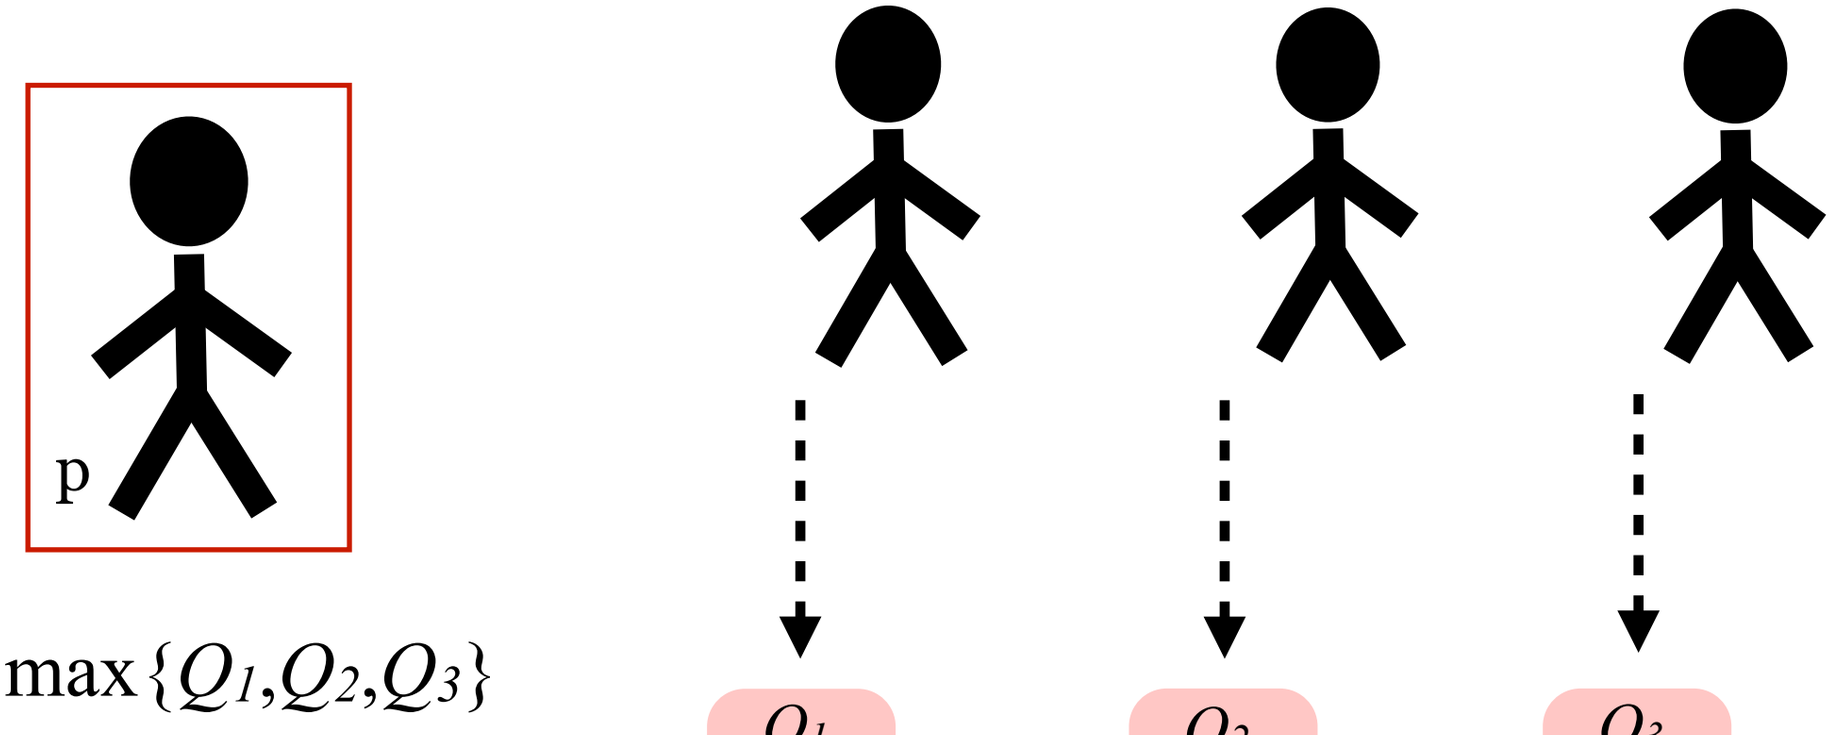

$u = \max\{Q_1, Q_2, Q_3\}$

Social welfare = *u – agent 1's production cost*
– agent 2's production cost
– agent 3's production cost

Efficient Crowdsourcing Contests [CJ, 2012]: The model

- A <u>principal</u> with *private value* seeks production of a good.
- A set of <u>agents</u> can individually produce goods.
  - Production yields **uncertain quality**.
  - Agents can expend variable *privately observed* **effort**; more effort leads to higher expected quality.
  - Agents have varying *private* **skill**; higher skill leads to higher expected quality.

Efficient Crowdsourcing Contests [CJ, 2012]: The model

- A <u>principal</u> with *private value* seeks production of a good.
- A set of <u>agents</u> can individually produce goods.
  - Production yields **uncertain quality**.
  - Agents can expend variable *privately observed* **effort**; more effort leads to higher expected quality.
  - Agents have varying *private* **skill**; higher skill leads to higher expected quality.

Will mostly focus on "constant skill" case today.

- Principal has value v ($) for a good with maximum quality
- Agent i with skill s_i chooses effort δ_i (which costs $\$\delta_i$)
  - a good is produced with quality distributed in a way that depends on v, s_i, and δ_i

- Principal has value v ($) for a good with maximum quality
- Agent i with skill s_i chooses effort δ_i (which costs $\$\delta_i$)
  - a good is produced with quality distributed in a way that depends on v, s_i, and δ_i

<u>Example</u>: quality Q_i uniformly distributed between 0 and $s_i\,\delta_i\,v$

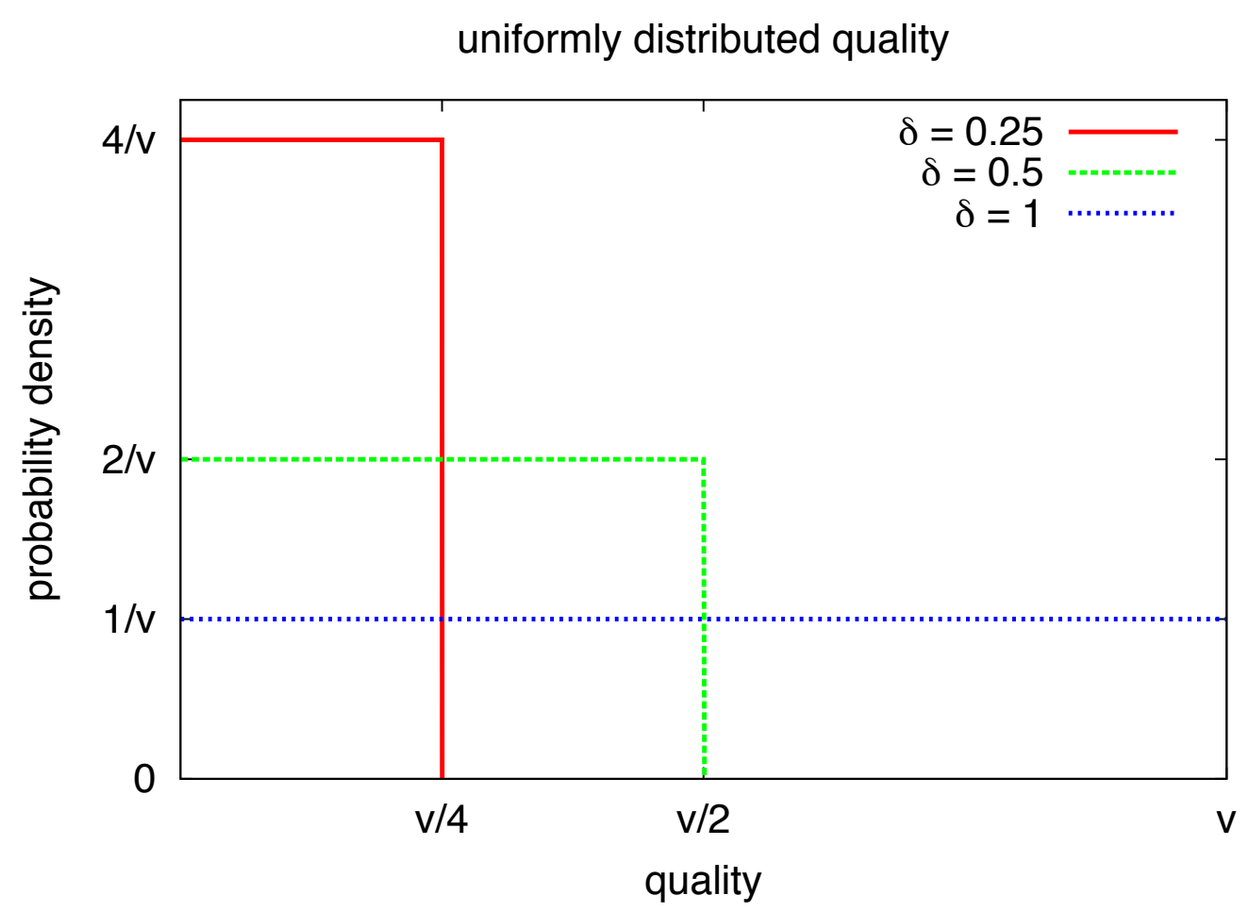

- Principal has value v ($) for a good with maximum quality
- Agent i with skill s_i chooses effort δ_i (which costs $\$\delta_i$)
  - a good is produced with quality distributed in a way that depends on v, s_i, and δ_i

<u>Example</u>: quality Q_i distributed normal with mean $s_i \delta_i v$

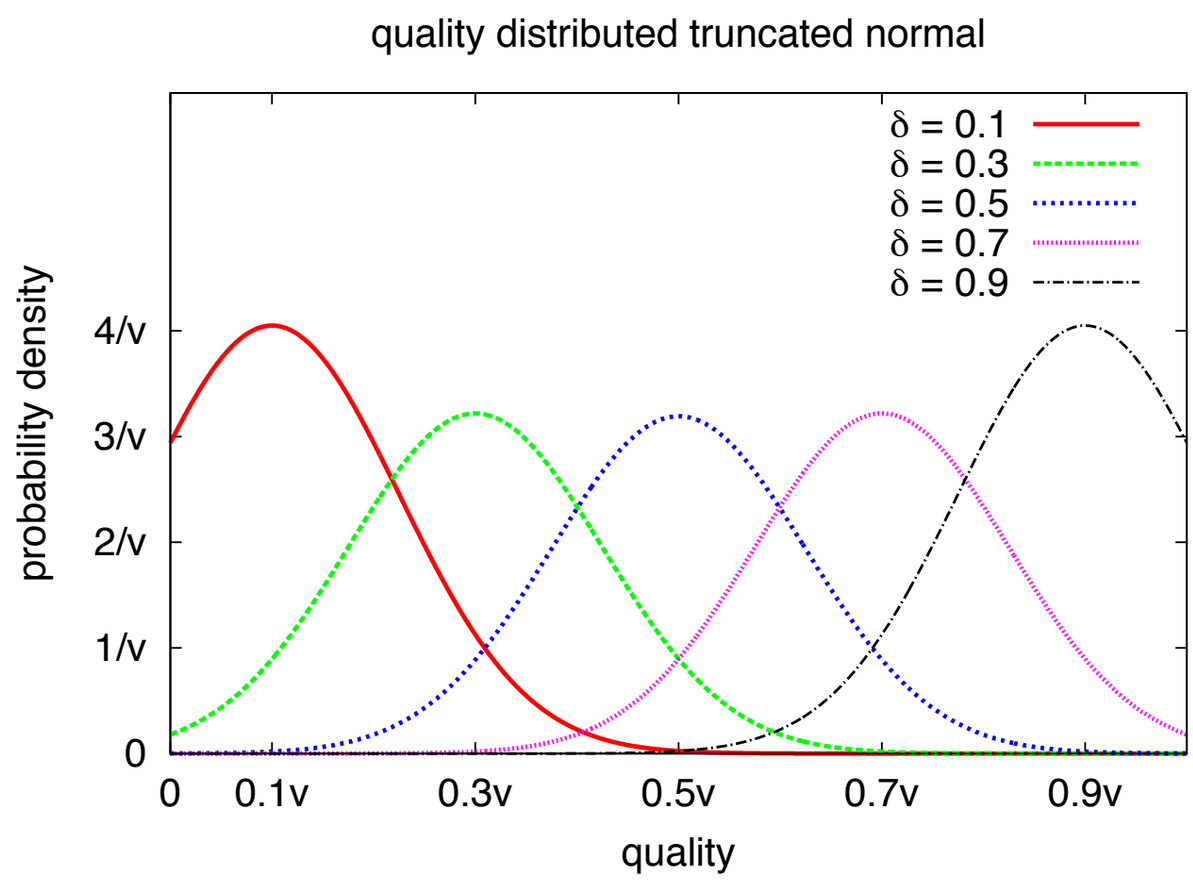

Seek to implement <u>efficient effort policy</u>, maximizing principal's obtained value *minus* sum of agents' costs (effort).

$$\mathbb{E}[\max_{i \in I} \mathcal{Q}_i(v, s_i, \delta_i)] - \sum_{i \in I} \delta_i$$

- Principal has value
- Agent i with skill s_i chooses effort δ_i (which costs $\$\delta_i$)
  - a good is produced with quality distributed in a way that depends on v, s_i, and δ_i

<u>Example</u>: quality Q_i distributed normal with mean $s_i \delta_i v$

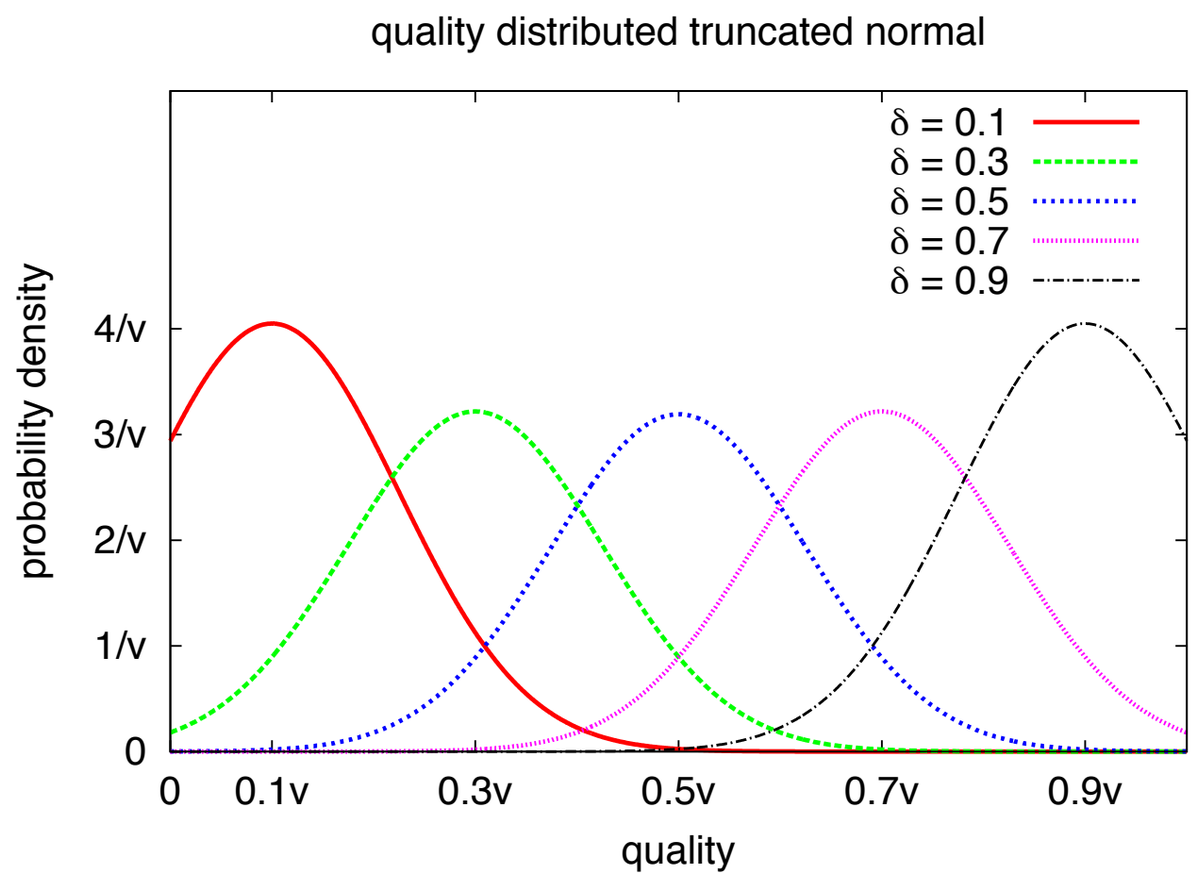

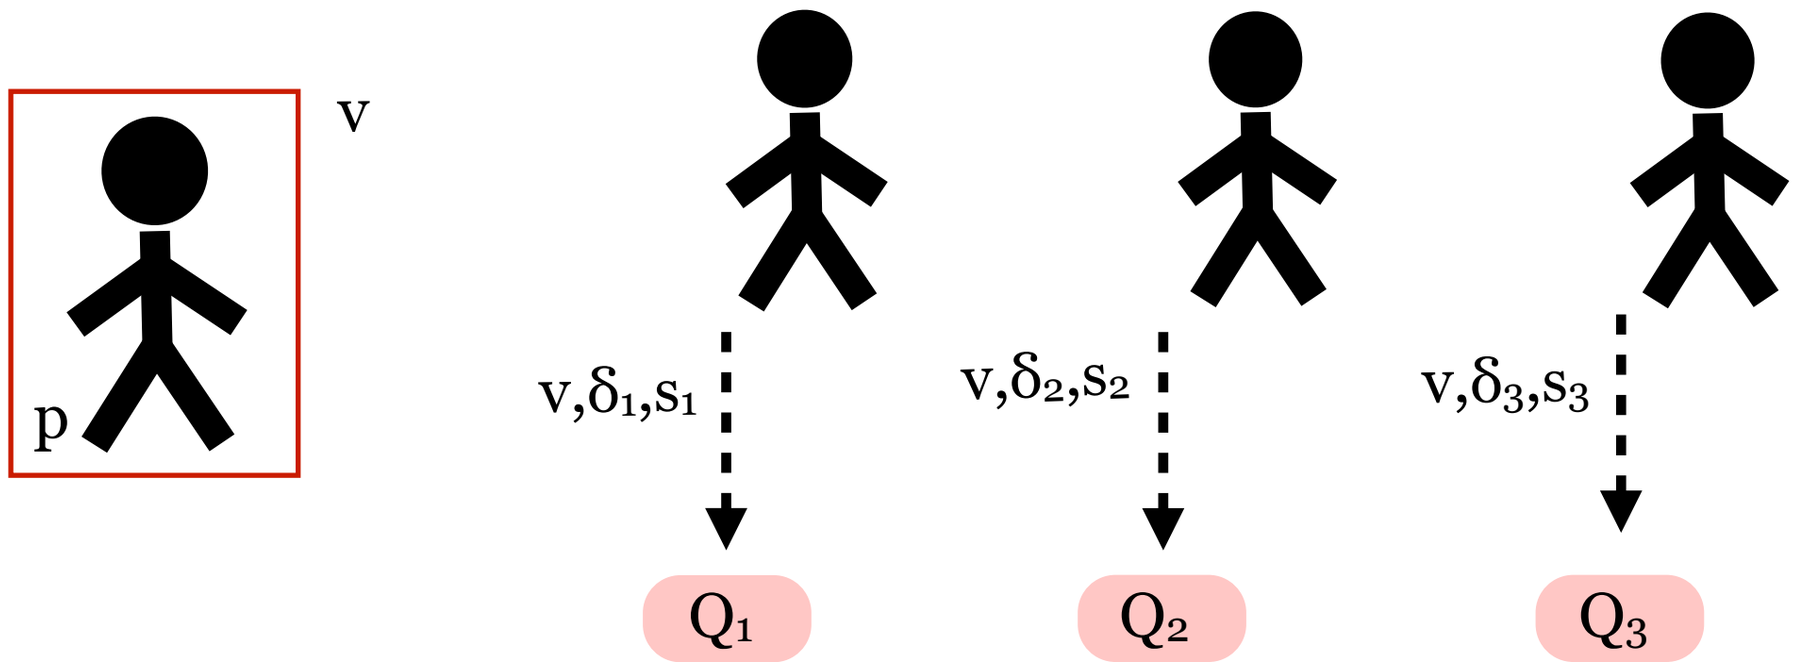

- *Quality* Q_i – dollar value to the principal of good that i produces – is a stochastic function of v, δ_i, and s_i.
- Social welfare equals: $\max\{Q_1,Q_2,Q_3\} - \delta_1 - \delta_2 - \delta_3$
- But since $\boldsymbol{v}$ and $\boldsymbol{s_i}$ are private, and $\boldsymbol{\delta_i}$ are privately observed, we need to incentivize principal and agents.

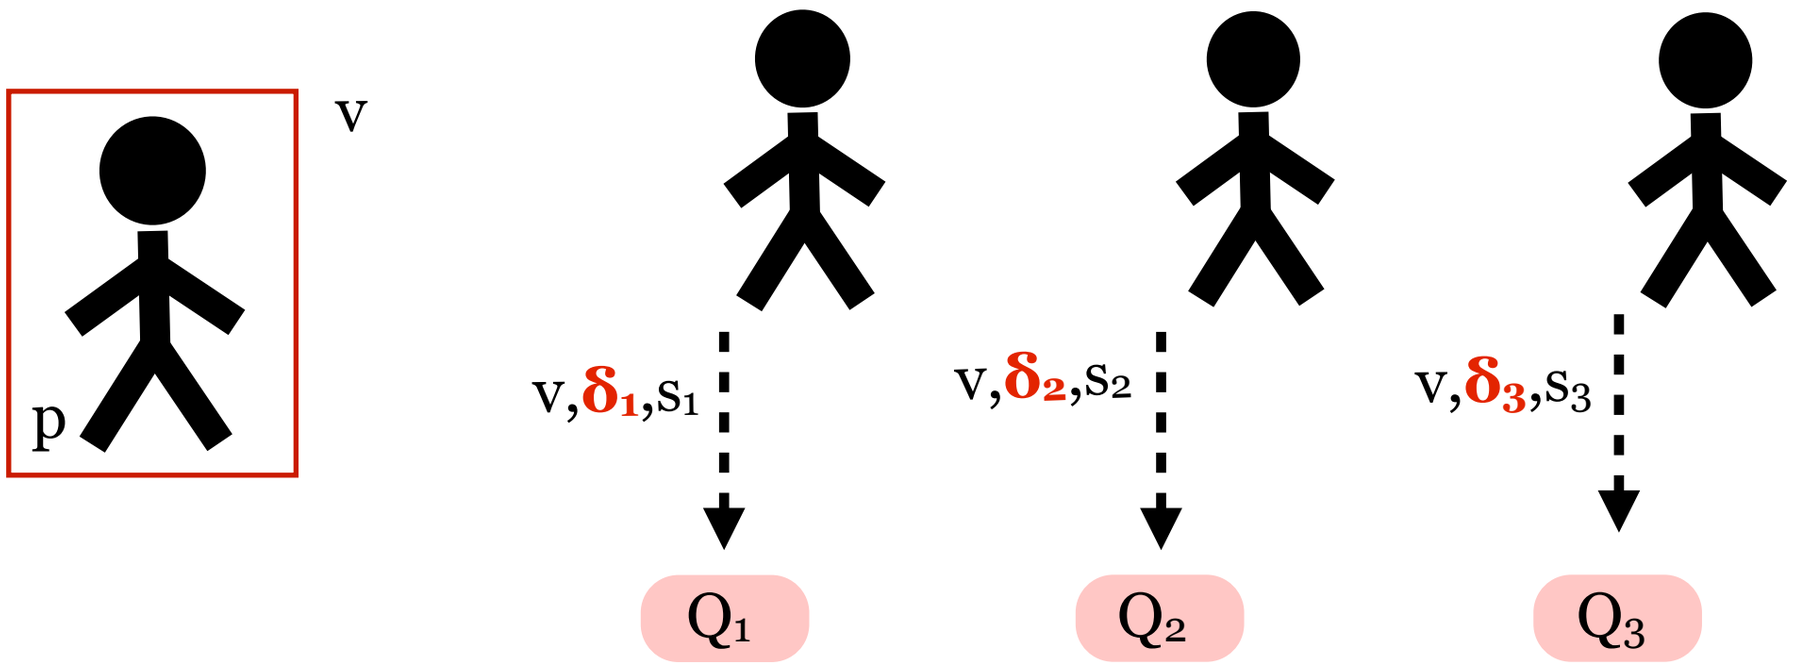

Efficient crowdsourcing involves:

1. A computational component:
   - Determine an effort policy that is efficient, i.e., maximizes sum of utilities (principal and agents).
2. An incentive component:
   - A payment mechanism that brings execution of such a policy into equilibrium.

Efficient effort policy

- In many cases, *extreme-effort* policies are optimal: each agent exerts either 0 effort or maximal effort.
- If extreme-effort policy is efficient, then determining efficient policy reduces to choosing *number* of participants.

Uniformly distributed quality

Theorem. For the constant skill, uniformly distributed quality case, a mechanism that elicits maximum-effort participation by m^* agents (and 0-effort participation by others) is efficient, where:

$$m^* = \begin{cases} \lfloor \sqrt{v} \rfloor - 1 & \textit{if } \lfloor \sqrt{v} \rfloor^2 + \lfloor \sqrt{v} \rfloor > v \\ \lfloor \sqrt{v} \rfloor & \textit{otherwise} \end{cases}$$

Uniformly distributed quality

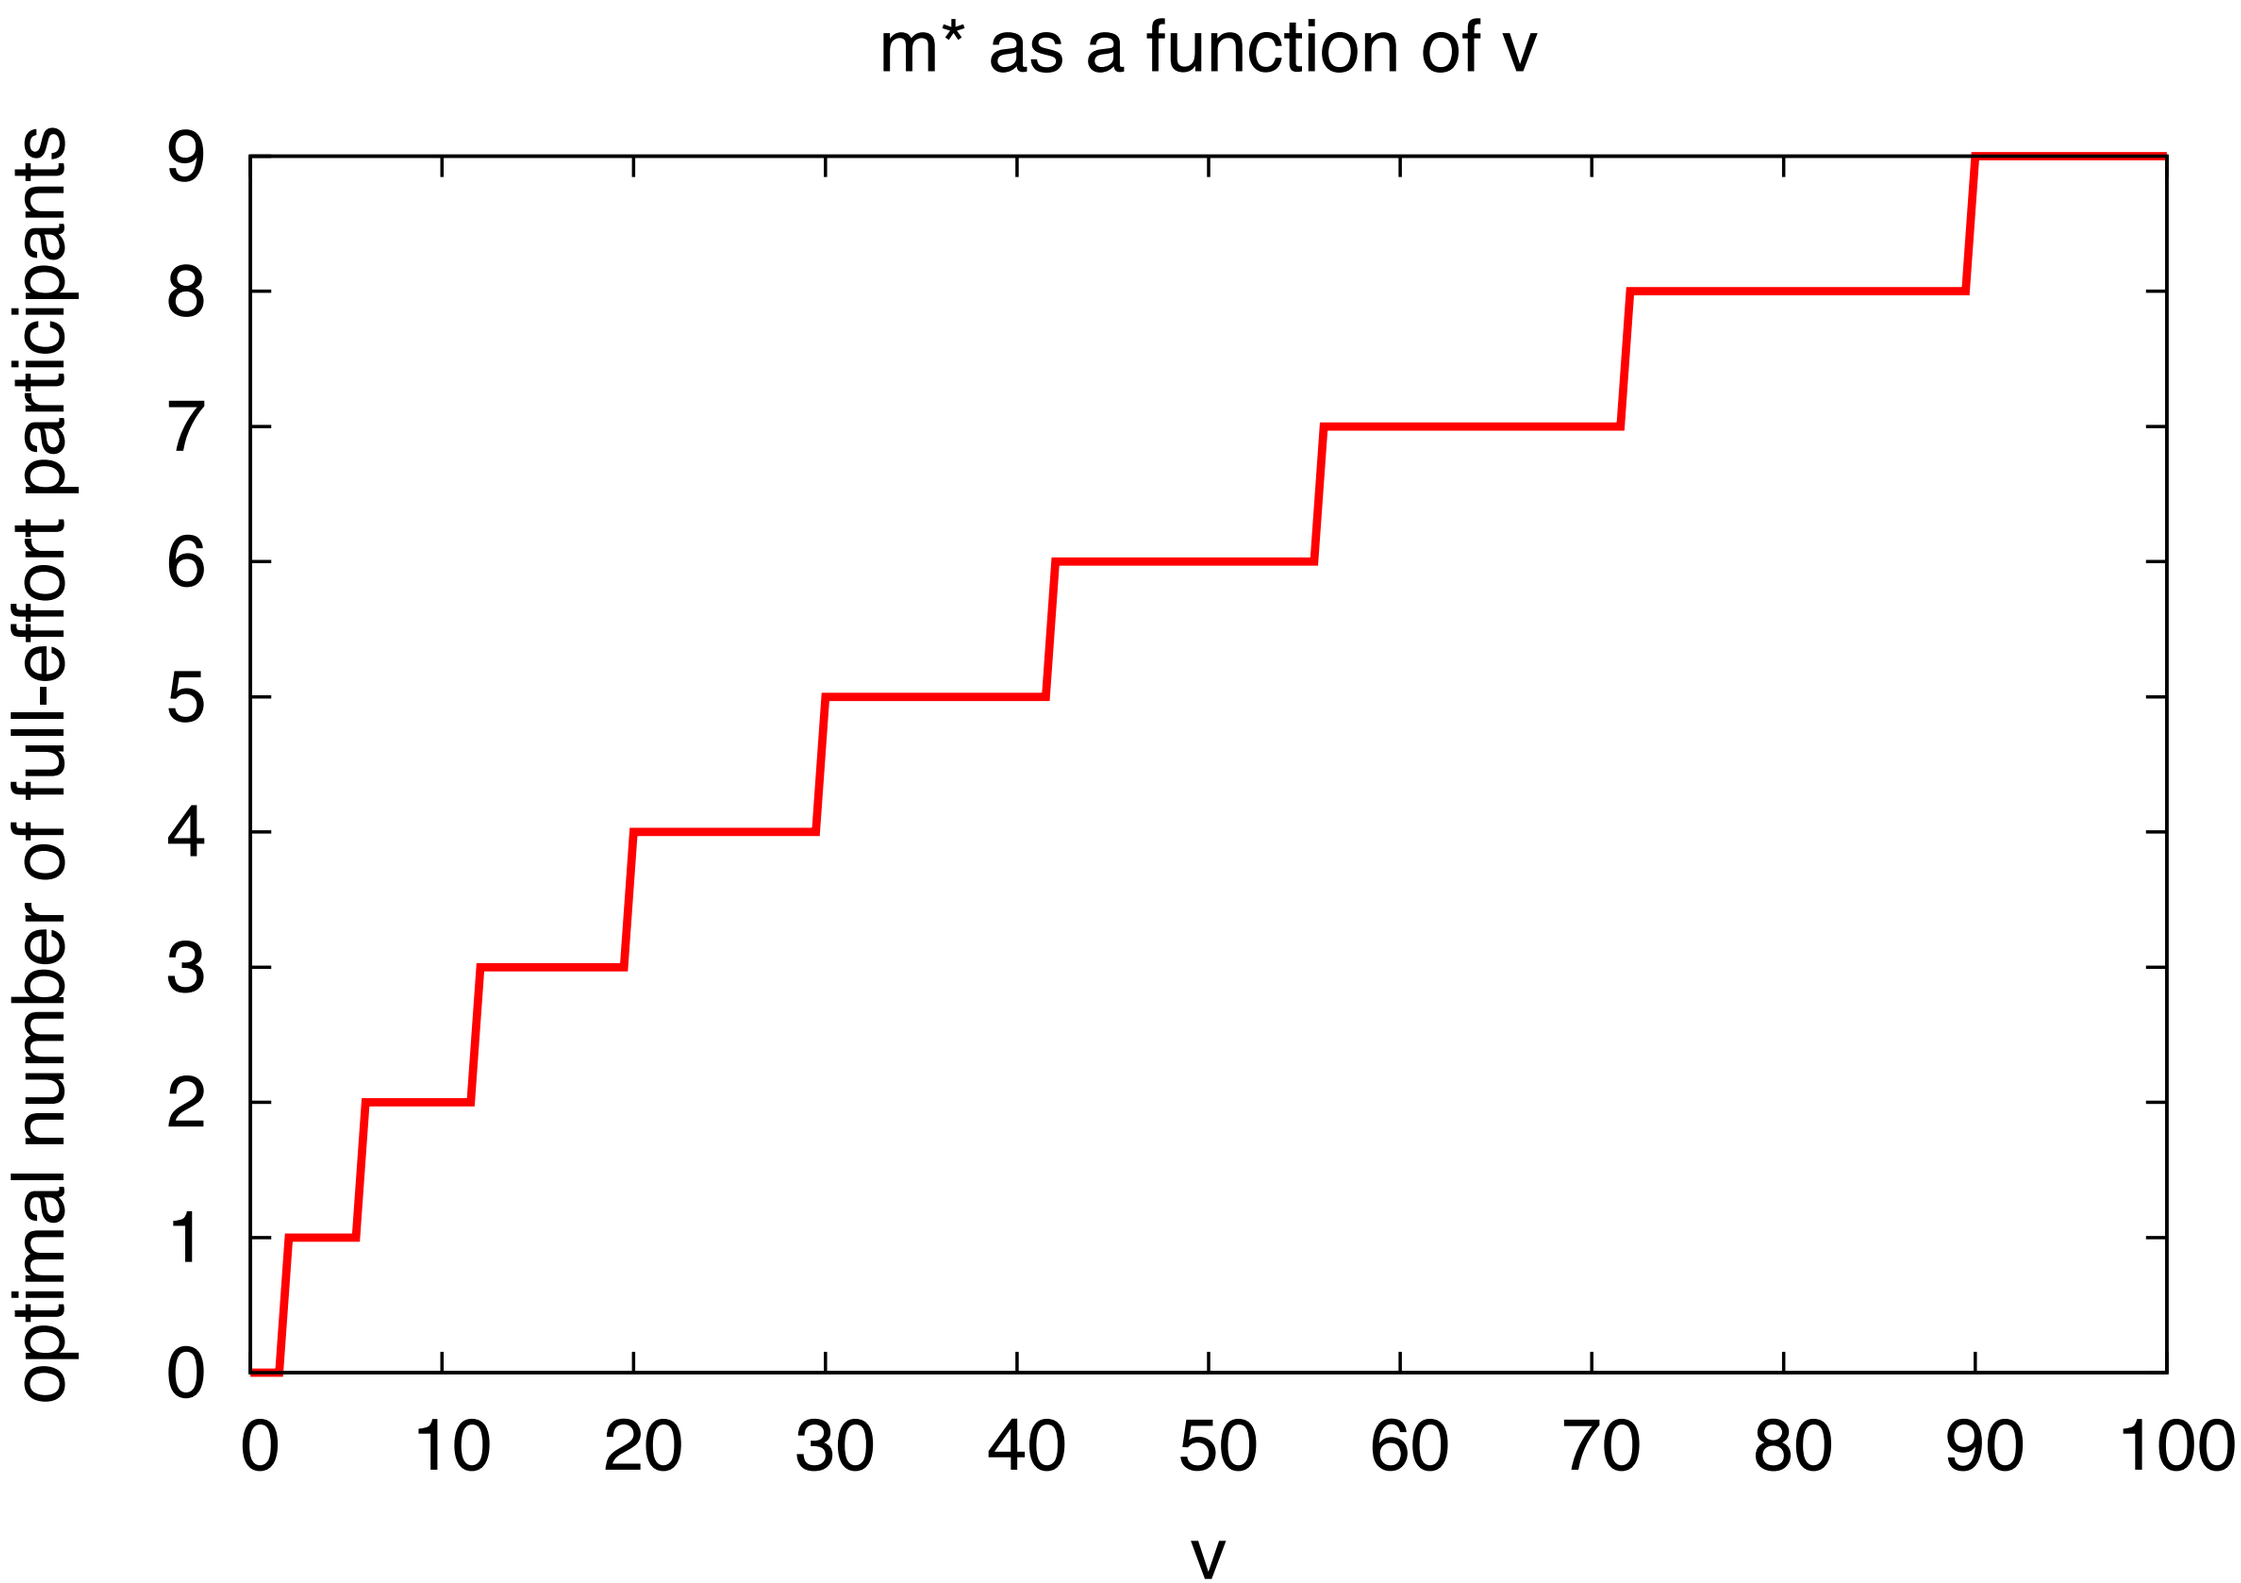

Normally distributed quality $\mu=\delta_i v$, $\sigma=v/8$

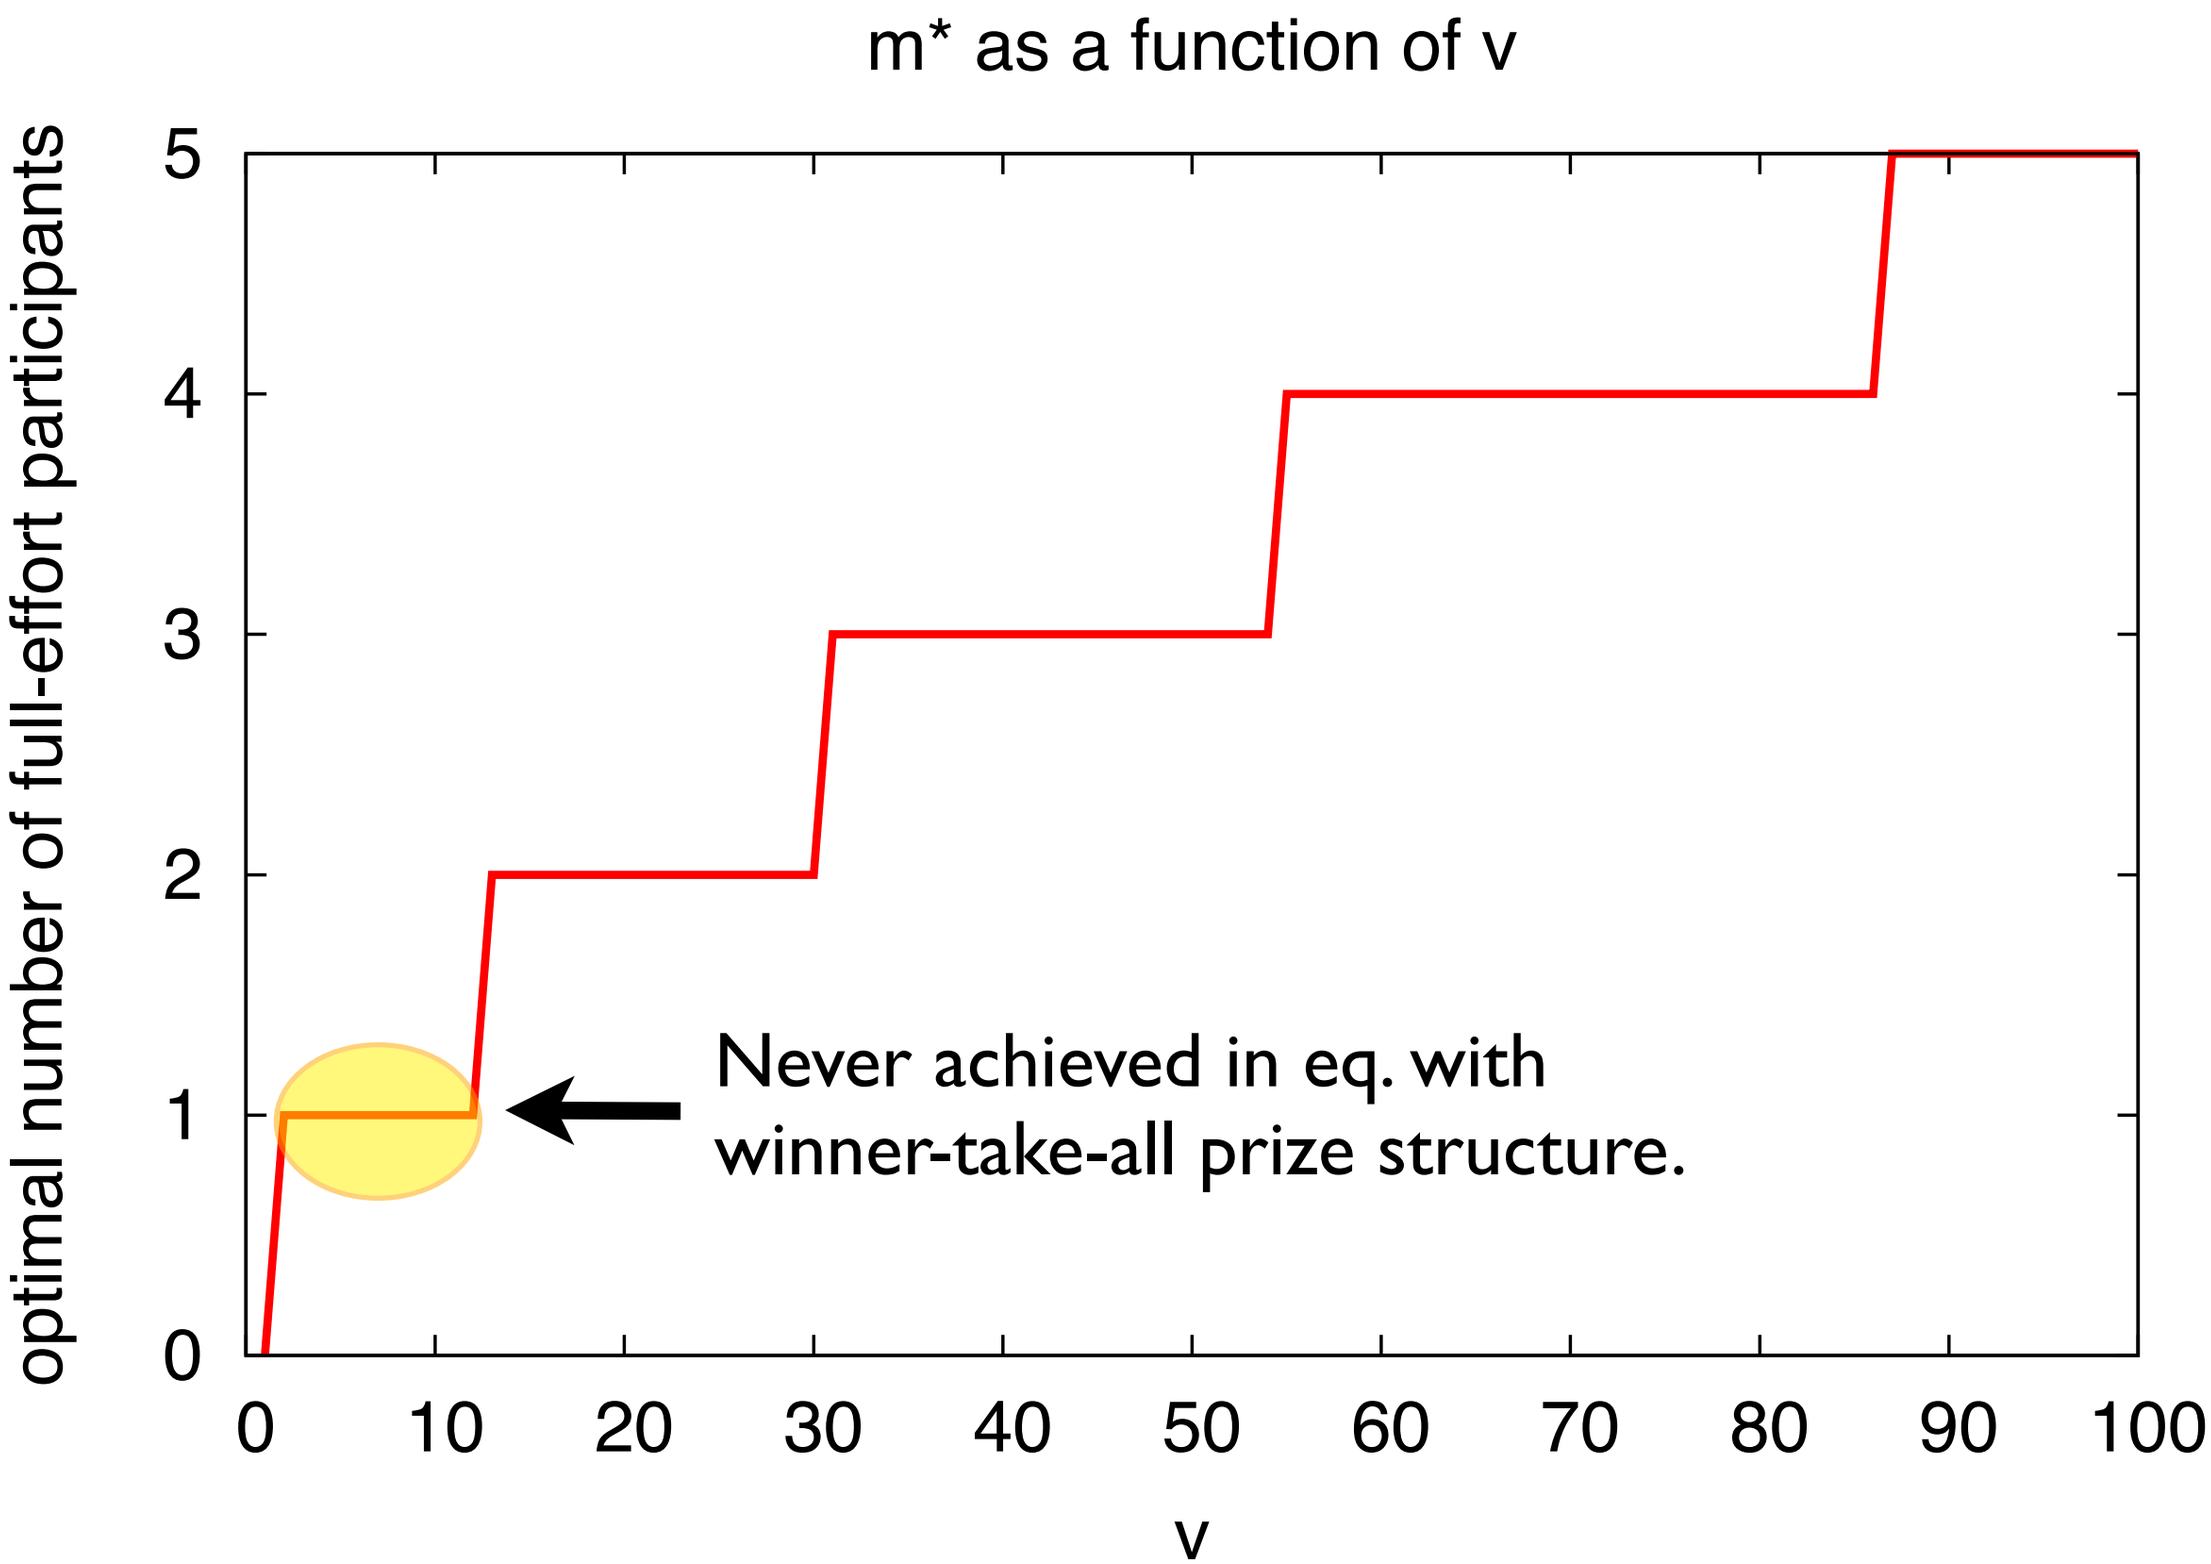

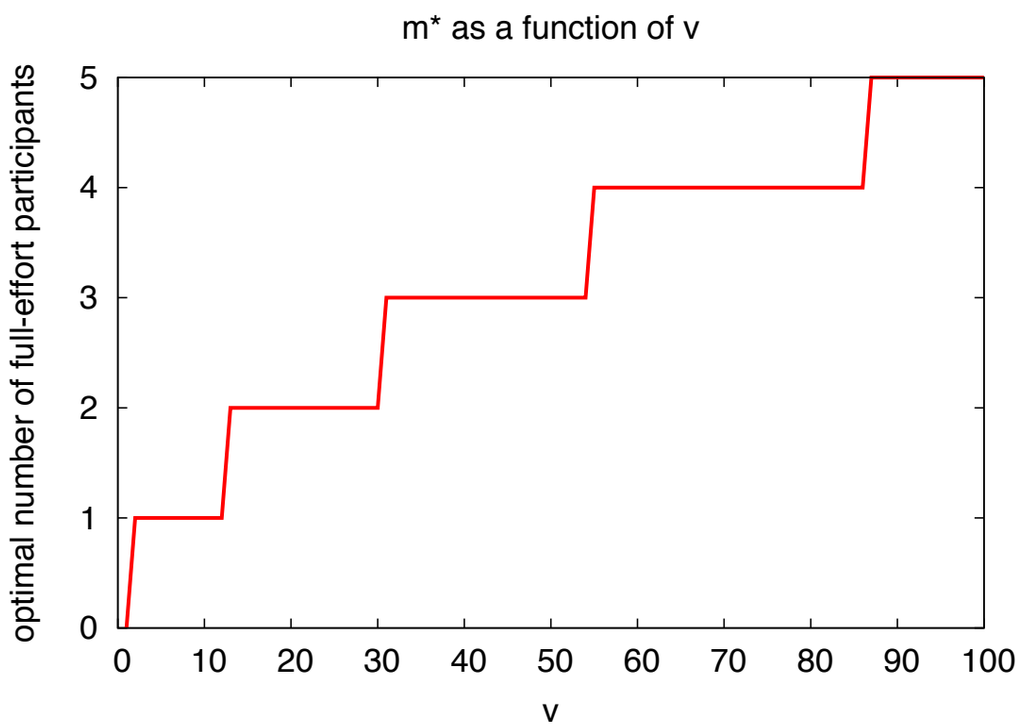

This is what we want to achieve. But can we?

Now for the incentives

Mechanism design

- The study of how to engineer incentives leading to desirable outcomes despite agent selfishness plus private information and/or autonomy.
- A few examples: auctions for allocating scarce resources; taxation to achieve a desired level of consumption; commissions to achieve sales performance.

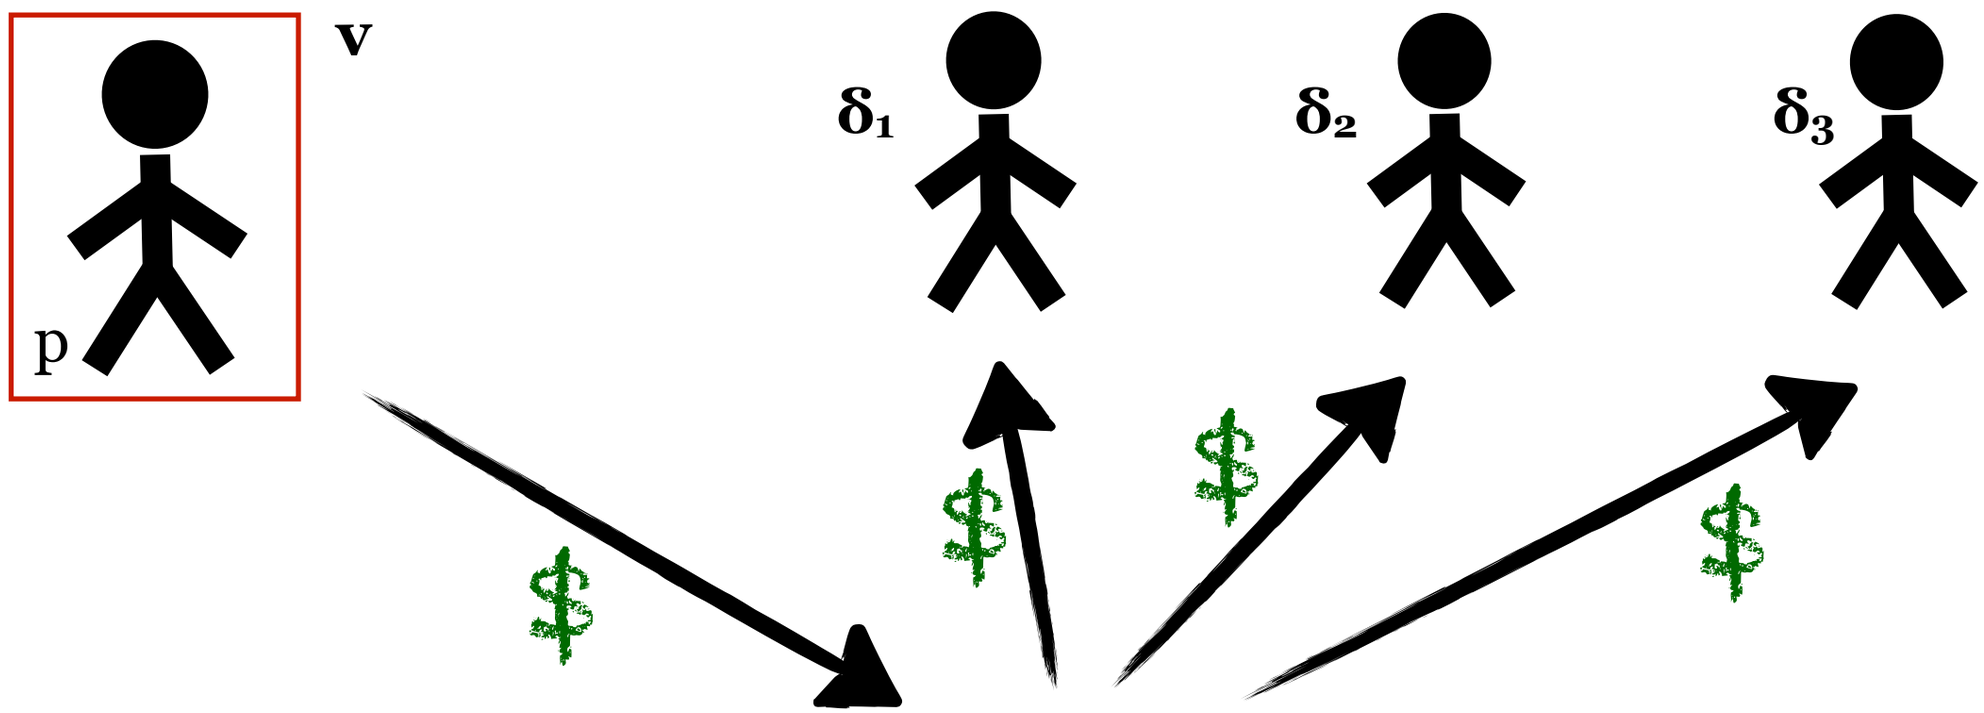

Make payments to get principal to report true value, and agents to exert prescribed amount of effort.

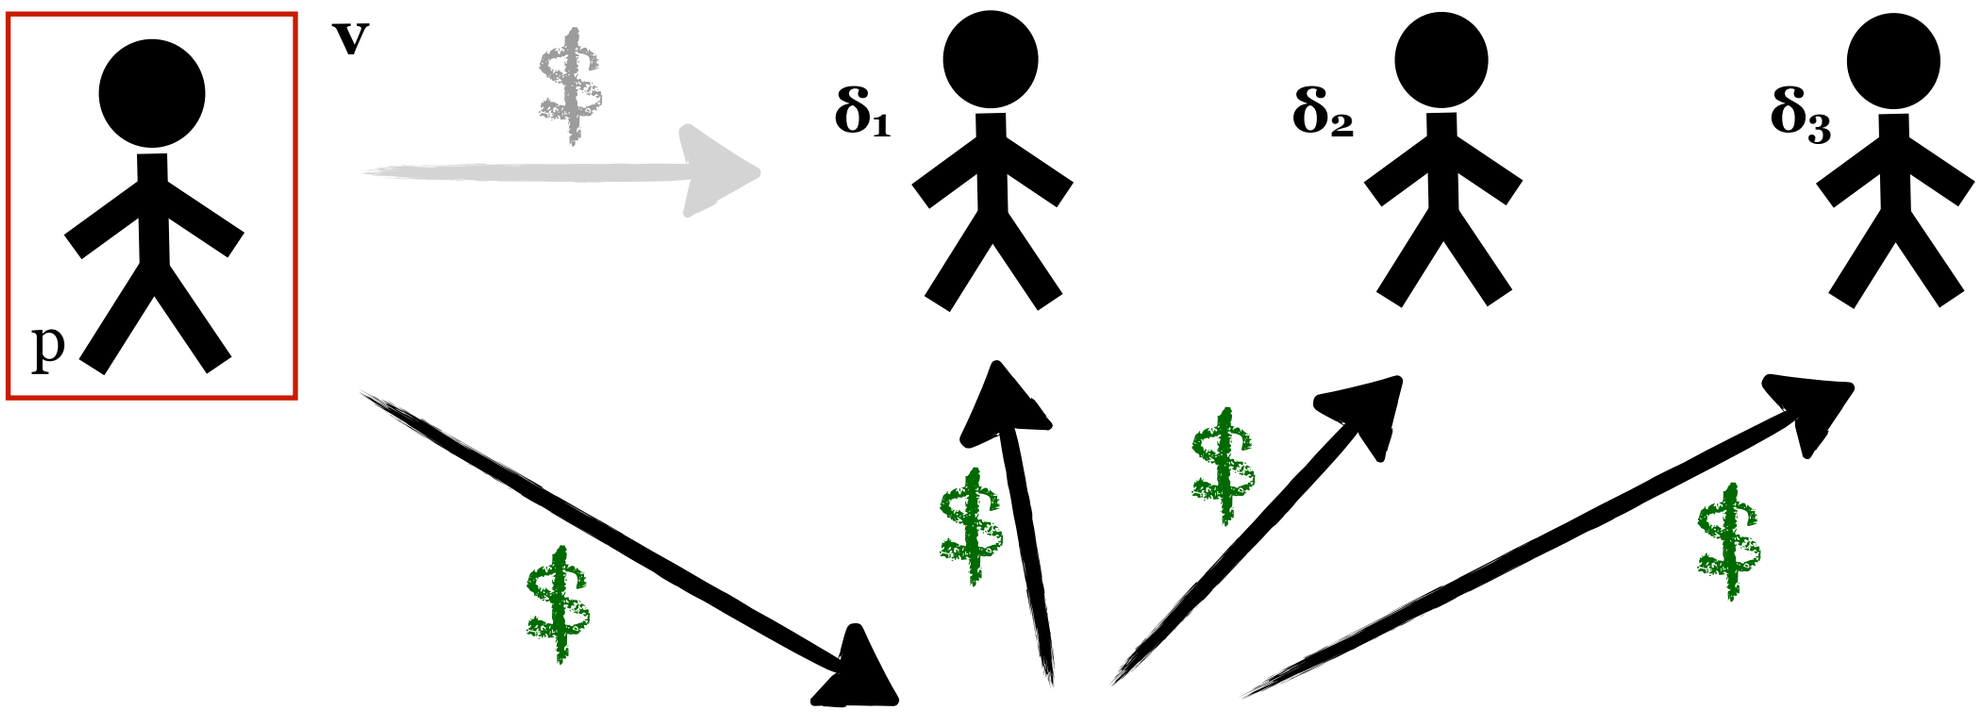

Make payments to get principal to report true value, and agents to exert prescribed amount of effort.

Goal: a mechanism that is...

- *Incentive compatible*: no one can benefit from deviating from honest participation
- *Individually rational*: everyone expects non-negative utility from participating honestly
- *No-deficit*: the mechanism cannot make positive aggregate payments

Mechanism for constant skill setting

- The principal reports v.
- Efficient effort levels $\delta_1, \ldots, \delta_n$ are computed.
- Each agent i is instructed to expend effort δ_i, and goods are produced with quality levels $Q_1, \ldots, Q_n$.

Mechanism for constant skill setting: payments

- The principal is charged: agents' aggregate prescribed effort ($\delta_1 + \delta_2 \ldots + \delta_n$).
- Each agent is paid:

  prescribed effort level +

  (highest quality level produced overall –

  highest quality level produced by *other* agents)
- Each agent is charged:

  E[highest quality level overall – highest quality level produced by other agents]

Theorem. This mechanism is efficient, incentive compatible, individually rational, and no-deficit in expectation for constant skill settings.

Mechanism for constant skill setting: **incentives**

- The principal is charged: agents' aggregate prescribed effort.
- Each agent is paid:

  prescribed effort level +

  (highest quality level produced overall –

  highest quality level produced by *other* agents)
- Each agent is charged:

  E[highest quality level overall – highest quality level produced by other agents]

has no effect on incentives

Principal's utility is:

highest quality level produced – aggregate effort expended

- The principal is charged: agents' aggregate prescribed effort.
- Each agent is paid:

  prescribed effort level +

  (highest quality level produced overall –

has no effect on incentives

Agent's utility is (proportional to):

highest quality level produced – aggregate effort expended

Mechanism for constant skill setting: **budget**

- The principal is charged: agents' aggregate prescribed effort.
- Each agent is paid:

  prescribed effort level +
  (highest quality level produced overall –
  highest quality level produced by *other* agents)
- Each agent

  E[highes
  level pr

Through this point, budget deficit equals quality difference between top two submissions.

Mechanism for constant skill setting: **budget**

- The principal is charged: agents' aggregate prescribed effort.
- Each agent is paid:

  prescribed effort level +
  (highest quality level produced overall –
  highest quality level produced by *other* agents)
- Each agent is charged:

  E[highest quality level overall – highest quality level produced by other agents]

Now budget is balanced in expectation.

Example (uniformly distributed quality)

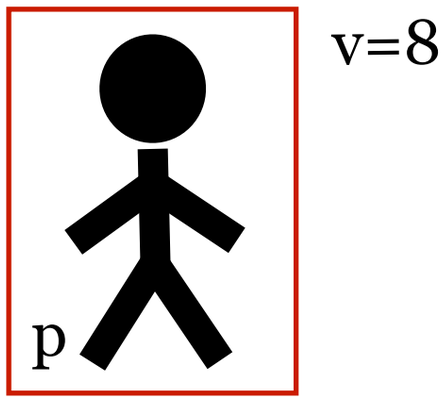

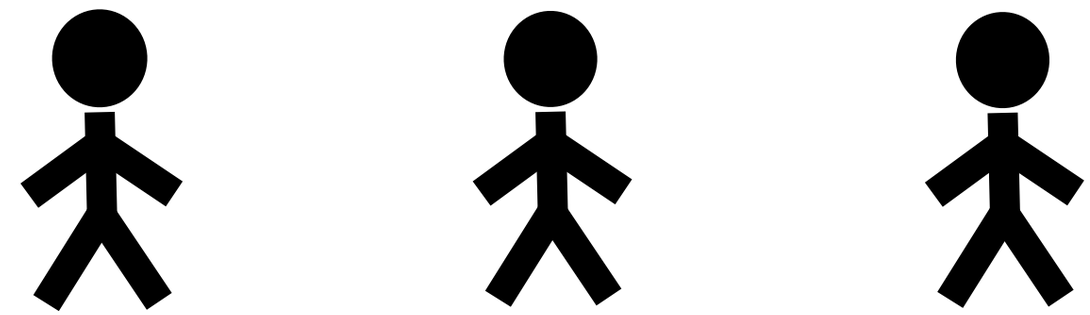

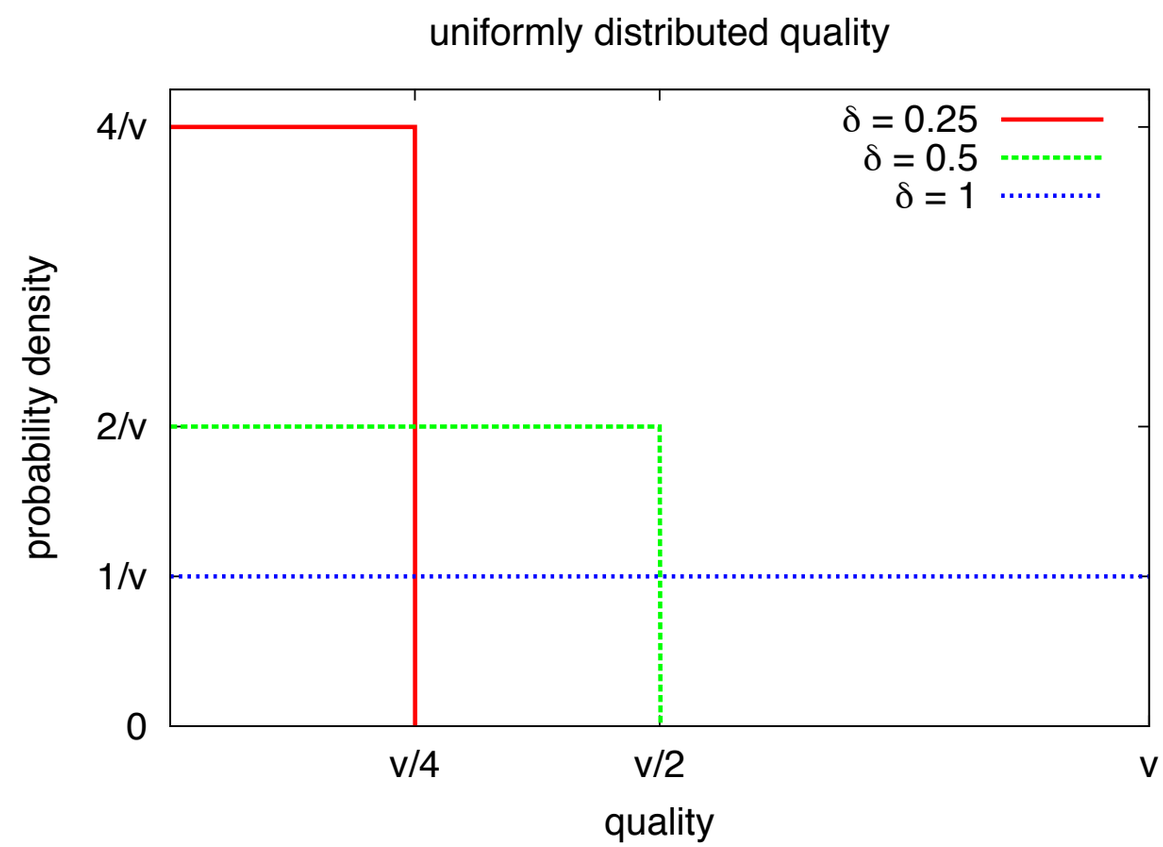

$$m^* = \begin{cases} \lfloor\sqrt{v}\rfloor - 1 & \textit{if } \lfloor\sqrt{v}\rfloor^2 + \lfloor\sqrt{v}\rfloor > v \\ \lfloor\sqrt{v}\rfloor & \textit{otherwise} \end{cases}$$

Example (uniformly distributed quality)

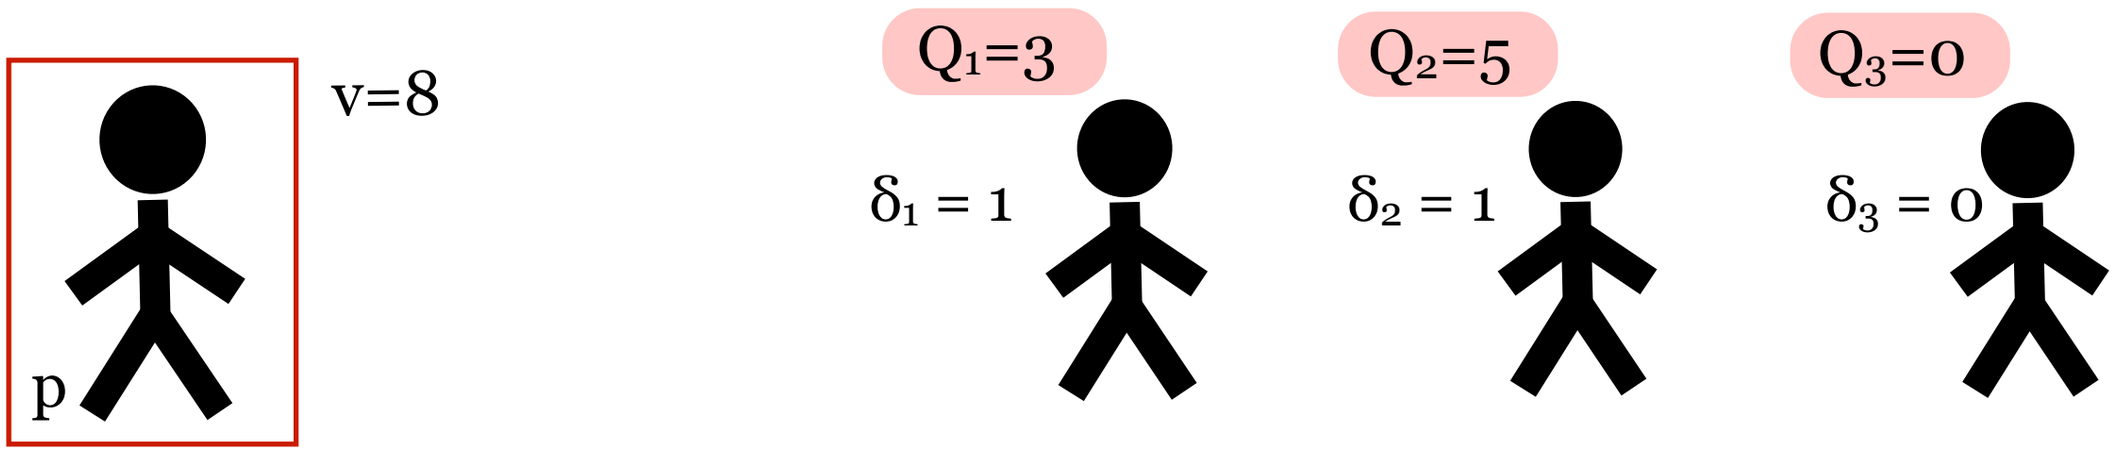

- Optimal policy has 2 agents exert full effort.

Example (uniformly distributed quality)

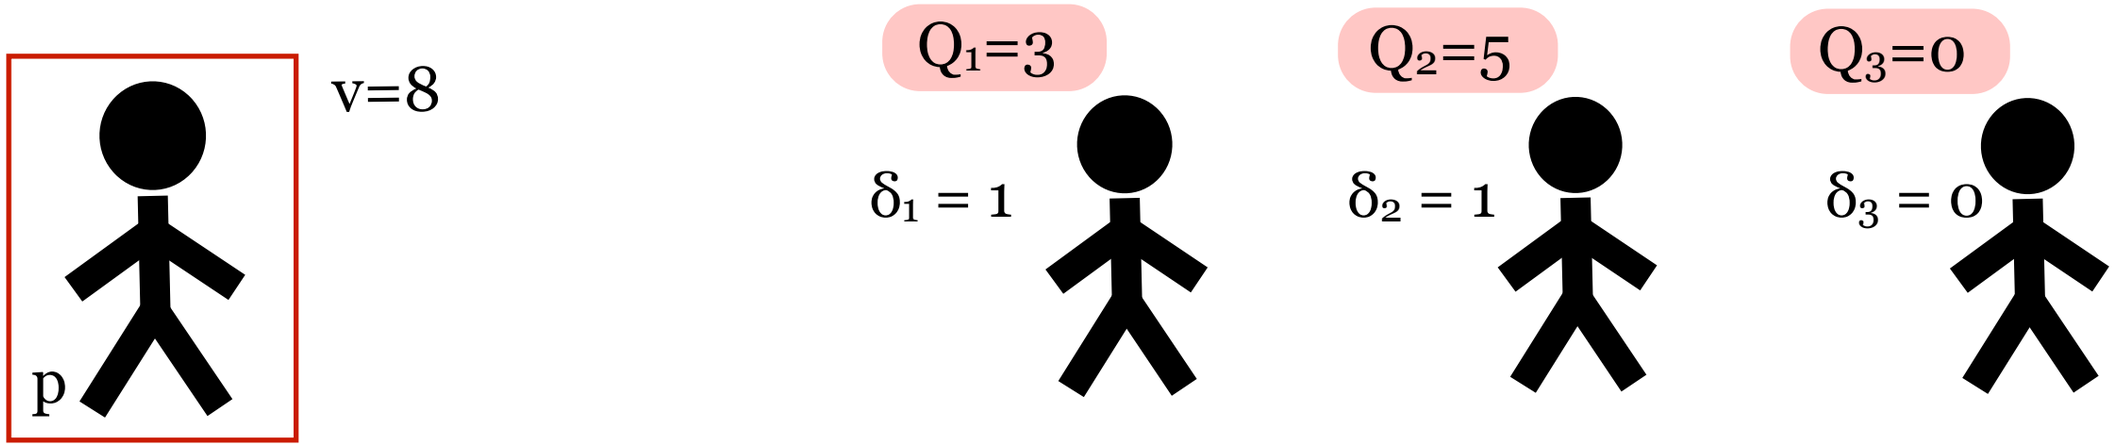

Must pay:
$\delta_1 + \delta_2 + \delta_3 = \mathbf{2}$

- Optimal policy has 2 agents exert full effort.

Example (uniformly distributed quality)

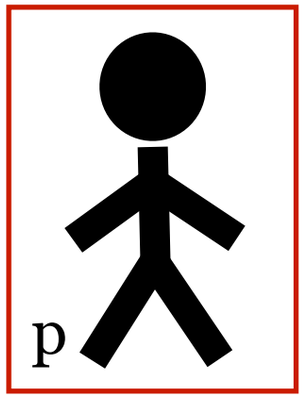

v=8

$Q_1=3$ $\delta_1 = 1$

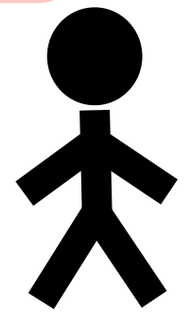

$Q_2=5$ $\delta_2 = 1$

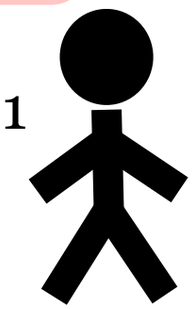

$Q_3=0$ $\delta_3 = 0$

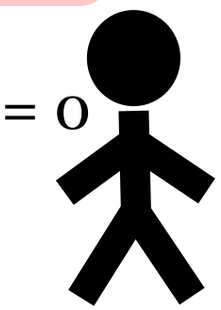

Must pay:
$\delta_1 + \delta_2 + \delta_3 = \mathbf{2}$

utility $= 5 - 2 = \mathbf{3}$

effort	$+1$	$+1$	$+0$
qual. diff	$+(5-5)$	$+(5-3)$	$+(5-5)$
E[qual. diff]	$-(16/3-4)$	$-(16/3-4)$	$-(16/3-16/3)$
is paid	$= \mathbf{-1/3}$	$= \mathbf{5/3}$	$= \mathbf{0}$
utility	$\mathbf{-4/3}$	$\mathbf{2/3}$	$\mathbf{0}$

Each's utility is non-negative *in expectation*, but not guaranteed.

Revenue $= 2 + 1/3 - 5/3 = \mathbf{2/3}$

In some cases, can charge principal *entry fee*, redistribute to agents to decrease odds of loss.

Private skill setting is more challenging.

Private skill

Theorem. There exists no mechanism that is efficient, incentive compatible, individually rational, and no-deficit in expectation.

However:

- In some cases weaker individual rationality concept may suffice.
  - Require commitment prior to revealing nature of task.

→ Efficient mechanism that is incentive compatible, no-deficit, and satisfies this notion of individual rationality.

public knowledge:

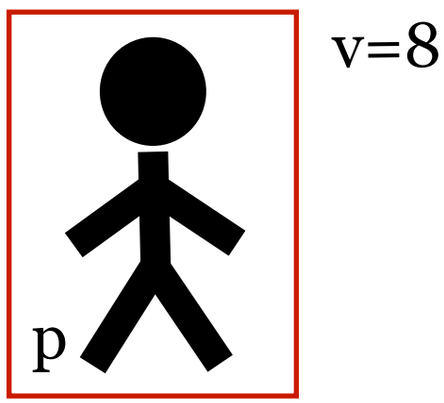

will participate regardless of v

$s_1 \sim U(0,1)$ $s_2 \sim U(0,1)$ $s_3 \sim U(0.5,1)$

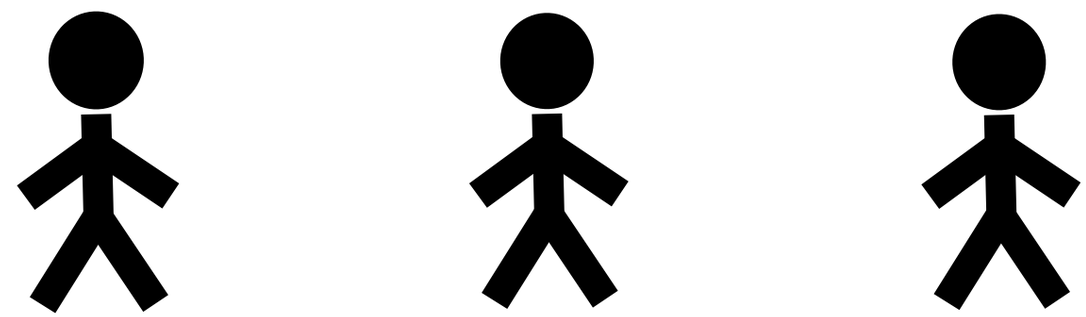

will "sign up" if forced to choose <u>now</u>

reveal task

private knowledge:

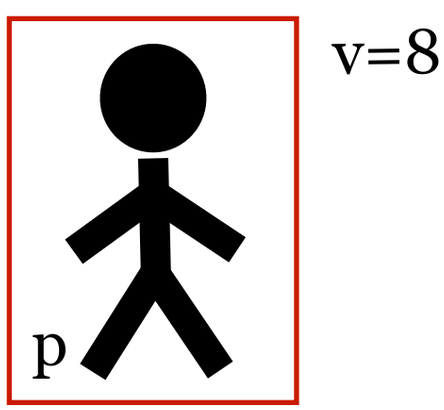

$s_1 = 0.7$ $s_2 = 0.2$ $s_3 = 0.5$

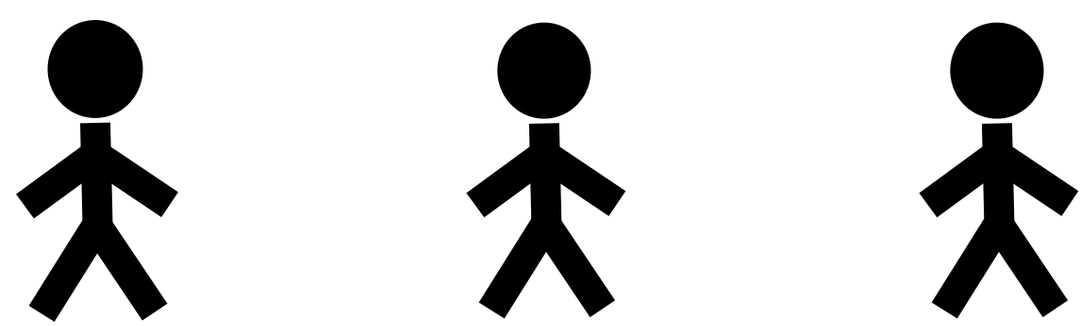

may or may not regret having participated

II. Mechanism for private skill setting

Let G equal minimum possible "expected value" for principal, given distribution over skill levels, and effort-to-quality distribution.

- The principal is charged: agents' aggregate prescribed effort, plus G.
- Each agent is paid: highest quality produced, minus effort prescribed for other agents, minus a balancing term independent of reported skill levels, plus $1/n$ times G.

Private skill mechanism

Theorem. The mechanism is incentive compatible, individually rational for the principal, individually rational for each agent *ex ante* of skill level realizations, and no-deficit in expectation *ex ante* of skill realizations.

Summary of [C&J, 2012]

+ Efficient mechanisms for crowdsourcing, very generally applicable
  + more efficiency → more attractive marketplace
– Lacks the simplicity of winner-take-all approach
  – in fact "contest" is now a misnomer
- Computing efficient policies can be hard (but can tractably handle lots of natural special cases)
- Important question: how big is the social welfare gain is in relevant cases?

High-level recap

- Long line of work in economics considering optimal contest design from perspective of **principal**, with recent contributions from [CHS, 2012] and [AS, 2009].
  - Significant assumptions that are questionable in typical web settings (deterministic production?).
- Then, [CJ, 2012] tries to take perspective of marketplace designer seeking to maximize **social welfare**, attracting principals *and* agents.
  - Incentives analysis is very generally applicable.
  - Computing optimal policies in the general case is hard and deserves more attention.