Removing sampling bias in networked stochastic approximation

Vivek Borkar*

IIT Bombay

July 2014

*Joint work with Raaz Dwivedi

Supported by DST

NETWORKED STOCHASTIC APPROXIMATION

- Directed or undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

- $\mathcal{N}(i) \subset \mathcal{V}$ is the set of neighbors of $i \in \mathcal{V}$

- each node $i \approx$ a computing element that performs the following iteration:

$$
\begin{aligned}
x_i(n+1) \;=\; & x_i(n) + a(n)[\sum_{j \in \mathcal{N}(i)} \xi_{ij}(n)\{h_{ij}(x(n), Y(n)) \\
& + M_i(n+1)\})].
\end{aligned}
$$

'Usual conditions': For $n \geq 0$ and

$$\mathcal{F}_n := \sigma(x(m), M(m), Y(m), \xi_{ij}(m), m \leq n, i, j \in \mathcal{V}),$$

- $\{Y(n)\}$ is a process taking values in a finite state space S and satisfying:

  $$P(Y(n+1) = j|\mathcal{F}_n) = p_{x(n)}(j|Y(n)), \; j \in S, n \geq 0,$$

  for a parametrized family of transition probabilities $\{p_x(\cdot|\cdot)\}, x \in \mathcal{R}^d$, on S such that the corresponding stochastic matrix P_x is irreducible and Lipschitz in x (the *Markov* noise),

- $\{M(n)\}$ is a square-integrable sequence adapted to $\{\mathcal{F}_n\}$ satisfying for $n \geq 0$,

  $E[M(n+1)|\mathcal{F}_n] = 0,\ E[\|M(n+1)\|^2|\mathcal{F}_n] \leq K(1+\|x(n)\|^2),$

  for some $K > 0$ (the *Martingale* noise),

- $a(n) > 0$ satisfy $\sum_n a(n) = \infty,\ \sum_n a(n)^2 < \infty$

- $h_{ij}(\cdot,\cdot) : \mathcal{R}^d \times S \mapsto \mathcal{R}$ Lipschitz in the first argument,

- $\{\xi_{ij}(n)\}$ independent $\{0,1\}$-valued random variables, $\xi_{ij}(n) = 1 \Longleftrightarrow i$ polls $j \in \mathcal{N}(i)$ at time n.

Notation:

- $\pi_x :=$ the unique invariant distribution under P_x

- $\hat{h}_{ij}(x) := \sum_k \pi_x(k) h_{ij}(x, k)$.

Assume 'stability': $\sup_n \|x(n)\| < \infty$ w.p.1.

Compare with the classical 'Robbins-Monro' scheme

$$x(n+1) = x(n) + a(n)[h(x(n)) + M(n+1)].$$

Tracks w.p.1 the asymptotic behavior of the o.d.e.

$$\dot{x}(t) = h(x(t)).$$

Our scheme tracks w.p.1 the asymptotic behavior of the o.d.e.

$$\dot{x}_i(t) = \sum_{j \in \mathcal{N}(i)} \lambda_{ij}(t) \widehat{h}_{ij}(x(t)),\ 1 \leq i \leq d,$$

where $\lambda_{ij}(t) \approx$ the 'instantaneous relative frequencies' with which i polls j.

This can have different and possibly undesired asymptotic behavior.

MODIFICATION:

Define $\nu(i,j,n) := \sum_{m=0}^{n} \xi_{ij}(m), \; n \geq 0$ ('local clocks').

Assume that:

1. There exists $\delta > 0$ such that $\forall i$,

$$\liminf_{n \uparrow \infty} \frac{\nu(i,j,n)}{n} \geq \delta \text{ a.s.} \tag{1}$$

(i.e., all components are updated 'comparably often').

2. $\{a(n)\}$ satisfy, for $A(n) := \sum_{m=0}^{n} a(m), c \in (0,1)$,

$$\sup_n \frac{a(\lfloor yn \rfloor)}{a(n)} < \infty \quad \forall\, y \in (0,1), \tag{2}$$

$$\frac{A(\lfloor yn \rfloor)}{A(n)} \overset{n\uparrow\infty}{\longrightarrow} 1 \text{ uniformly in } y \in (c,1]. \tag{3}$$

These are satisfied, e.g., by $a(n) = \frac{1}{n}, \frac{1}{n \log(n)+1}$ etc., but not by, e.g., $\frac{1}{n^{\frac{2}{3}}}$.

($a(n) \downarrow$ 'fast enough' $\Longrightarrow A(n) \uparrow$ sufficiently slowly).

Replace our iteration by by

$$x_i(n+1) \;=\; x_i(n) + [\sum_{j\in\mathcal{N}(i)} a(\nu(i,j,n))\xi_{ij}(n)\{h_{ij}(x(n),Y(n)) \\ +M_i(n+1)\}].$$

If stable, then it tracks w.p.1 the asymptotic behavior of the o.d.e.

$$\dot{x}(t) = \frac{1}{d}h(x(t)),$$

i.e., $\lambda_{ij}(t) \equiv \frac{1}{d} \;\forall\; i,j \in \mathcal{N}(i), t > 0$. where d as before is the dimension of $x(t)$, equivalently, the number of nodes.

$\Longrightarrow$ the asymptotic behavior of this o.d.e. is the same as that of $\dot{x}(t) = h(x(t))$

(The two are time-scaled versions of each other – set $\tau := \frac{t}{d}$.)

$\Longrightarrow$ identical trajectories, only the speed with which they are traversed is affected.

Communication delays can also be handled.

A Reputation System (Truong et al)

- Experts $\{1, \cdots, d\}$ with ratings ('reputation') $p_t^i, t \geq 0, 1 \leq i \leq d$,

- equal initial reputation: $p_0^i = \frac{1}{d} \; \forall i$,

- $x_t^i \in [0, 1]$: expert i's predictions of i.i.d. observations $y_t \in \{0, 1\}$,

- $\widehat{y}_t$: weighted prediction given by

$$\widehat{y}_t := \frac{\sum_{i \in E_t} p_t^i x_t^i}{\sum_{i \in E_t} p_t^i},$$

- $E_t :=$ the set of experts active at time t,

- p_t^i according to

$$\begin{aligned} p_{t+1}^i &= p_t^i \frac{x_t^i}{\widehat{y}_t} \text{ if } i \in E_t,\ y_t = 1, \\ &= p_t^i \frac{1 - x_t^i}{1 - \widehat{y}_t} \text{ if } i \in E_t,\ y_t = 0, \\ &= p_t^i \text{ if } i \notin E_t. \end{aligned}$$

Assumption: The distribution of $I\{i \in E_t\}, t \geq 0$, is stationary and symmetric in i

$\Longrightarrow p_t^i \stackrel{t\uparrow\infty}{\rightarrow} 1$ w.p.1 for the best expert if unique, otherwise the scheme oscillates between best experts.

'*Assumption*' above necessitated by the sampling bias: algorithm favors experts who opine more often.

Alternative scheme:

$$
\begin{aligned}
p^i_{t+1} &= \Gamma(p^i_t[1 + a(\nu(i,t))I\{i \in E_t\}w^i_t \\
&\quad - \sum_j a(\nu(j,t))I\{j \in E_t\}p^j_t w^j_t)]),
\end{aligned}
$$

where,

- $w^j_t := y_t x^j_t + (1 - x^j_t)(1 - y_t)$,
- $\nu(j,t) := \sum_{m=0}^t I\{j \in E_m\}$,

- $\Gamma(\cdot)$ is the projection onto the d-dimensional probability simplex S,

- $y_t \in [0, 1]$,

- E_t to be i.i.d. (can be relaxed), but *not necessarily symmetric*.

Let $z_i := E[w_t^i], 1 \leq i \leq d$, and without loss of generality, let $z_1 > z_j, j \neq 1$ (i.e., expert 1 is the best expert).

Theorem $p_t^1 \stackrel{t\uparrow\infty}{\to} 1$ w.p.1.

Remarks:

- Analysis based on the limiting o.d.e., which is a simple instance of the 'replicator dynamics'.

- Use constant stepsize for tracking slowly varying environment.

- objective 'ordinal' $\Longrightarrow$ 'convergence' fast.

- Applications to networked control???

Numerical experiments

- We present numerical results for two different cases for the reputation system.

- We label the iterations by index $n \geq 0$ and thus n plays the role of t.

- For a given i, we generate $I(i \in E_n)$ in an i.i.d. fashion for different i.

- For projection Γ, in simulations we use a slight modification followed by normalization

$$\begin{aligned} \text{First calculate, } p_i(n+1) &= \max(\epsilon, Y), \ \forall i, \\ \text{And then normalize, } p_i(n+1) &= \frac{p_i(n+1)}{\sum_j p_j(n+1)}, \end{aligned} \quad (4)$$

where Y denotes the argument of Γ in RHS of alternative scheme with n replacing t and $\epsilon = 10^{-6}$ is a small number to prevent the algorithm from accidentally getting stuck at a lower dimensional face of the probability simplex (Note that these are invariant sets for the iteration).

- For clarity we depict the variation of p_i against iterations for only three experts with largest z_i values. Using the convention that an expert i is better than expert j if $z_i > z_j$, we call these three experts as 'Best 3 Experts'. If required, we refer to the unique best expert as expert i^*. We also define $\nu_i := E(I(i \in E_n))$ (note the close relation to $\nu(i,t)$).

- The convergence rate of $p_{i^*}(n)$ to 1 depends upon values of z_{i^*} and ν_{i^*} relative to other z_i's and ν_i's resp. and can be boosted by changing the step size schedule from $a(n) = \frac{1}{n}$ to $a(n) = \frac{1}{[n/K]+1}$ where [x] denotes greatest integer not greater than x and K is a suitably chosen large integer.

- We use the modified step schedule with $K = 100$ for both cases. This schedule continues to satisfy conditions given in modification, but has a slower decrease, leading to faster convergence at the expense of somewhat higher fluctuations. (This is a standard trade-off in stochastic approximation.)

- We provide two figures for each case, one of them depicting transience (and fluctutations because of the modification for faster convergence) and other showing the (steady state) convergence result.

Case 1

For a reputation system with 20 experts, we generate $x_i(n)$'s and $y(n)$ as independent random variables uniformly distributed in $[0, 1]$ with randomly pre-assigned means. Transience plot shows that $p_{i^*}(n)$ is far from 1 because i^* has not opined sufficient number of times to be identified as the best expert. While steady state plot shows that finally the iterates converge to Dirac measure 1, with value 1 for the expert with highest z_i, though ν_{i^*} is 'approximately half' of the second best expert.

Transient Behavior of Best 3 Experts' Reputation

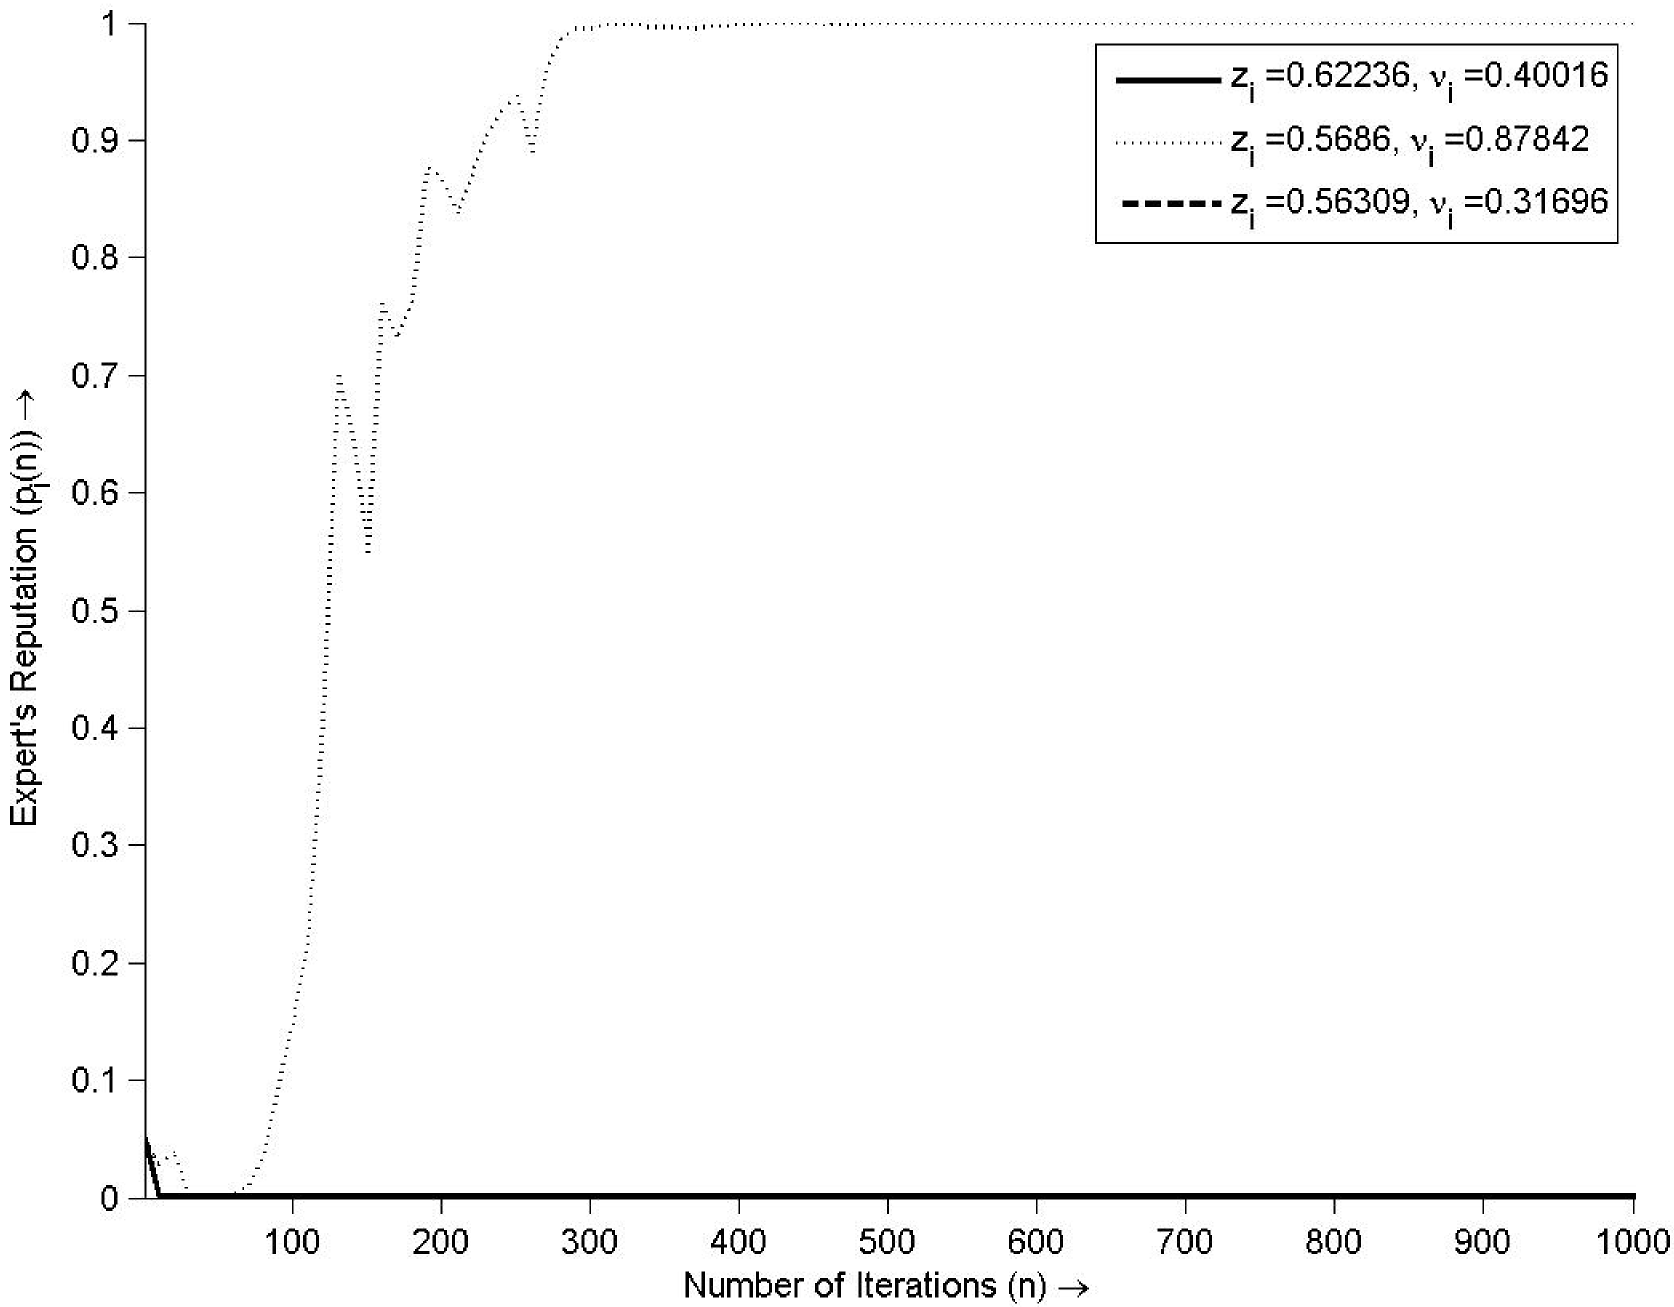

1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
100
200
300
400
500
600
700
800
900
1000
Expert's Reputation ($p_i(n)$) →
Number of Iterations (n) →
z_i =0.62236, ν_i =0.40016
z_i =0.5686, ν_i =0.87842
z_i =0.56309, ν_i =0.31696

Steady State Behavior of Best 3 Experts' Reputation

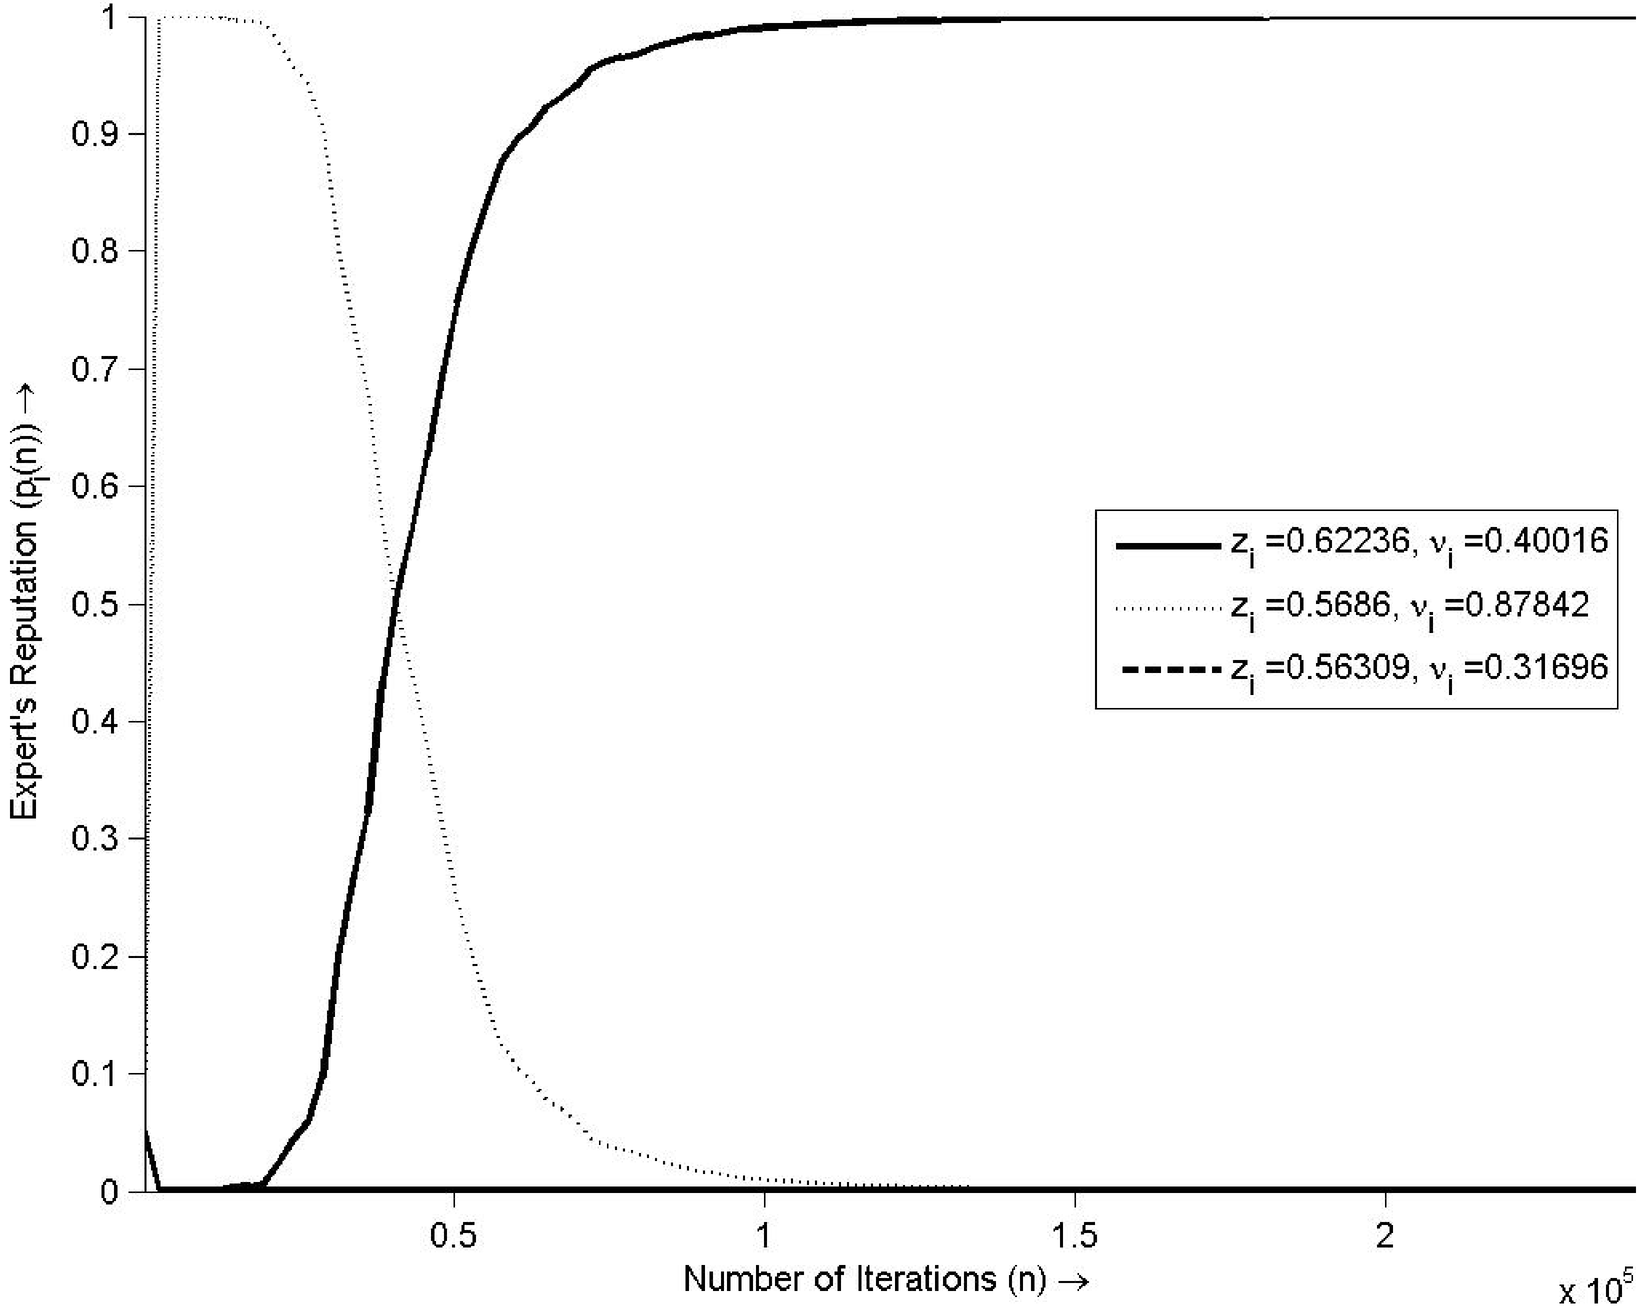

1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Expert's Reputation ($p_i(n)$) →
0.5
1
1.5
2
Number of Iterations (n) →
x 10^5
z_i =0.62236, ν_i =0.40016
z_i =0.5686, ν_i =0.87842
z_i =0.56309, ν_i =0.31696

Case 2

We simulate a reputation system with 10 experts. Here, we directly generate w_i's. The z_i values are pre-assigned deterministically with one best expert i^* such that $z_{i^*} = 2z_i, \ \forall i \neq i^*$.

However, the best expert is '10 times less likely' to opine than any other expert. That is, $\nu_{i^*} = \frac{1}{10}\nu_i, \forall i \neq i^*$. We assign such ratios to demonstrate that the algorithm is in fact successful in removing the sampling bias. Transience plot shows that initially there are great fluctuations but eventually $p_{i^*}(n)$ does converge to 1 as evident from Steady State Plot.

As compared to the previous case, the number of iterations for convergence are much larger because of the 'very rare' opining by the 'best expert' .

Transient Behavior of Best 3 Experts' Reputation

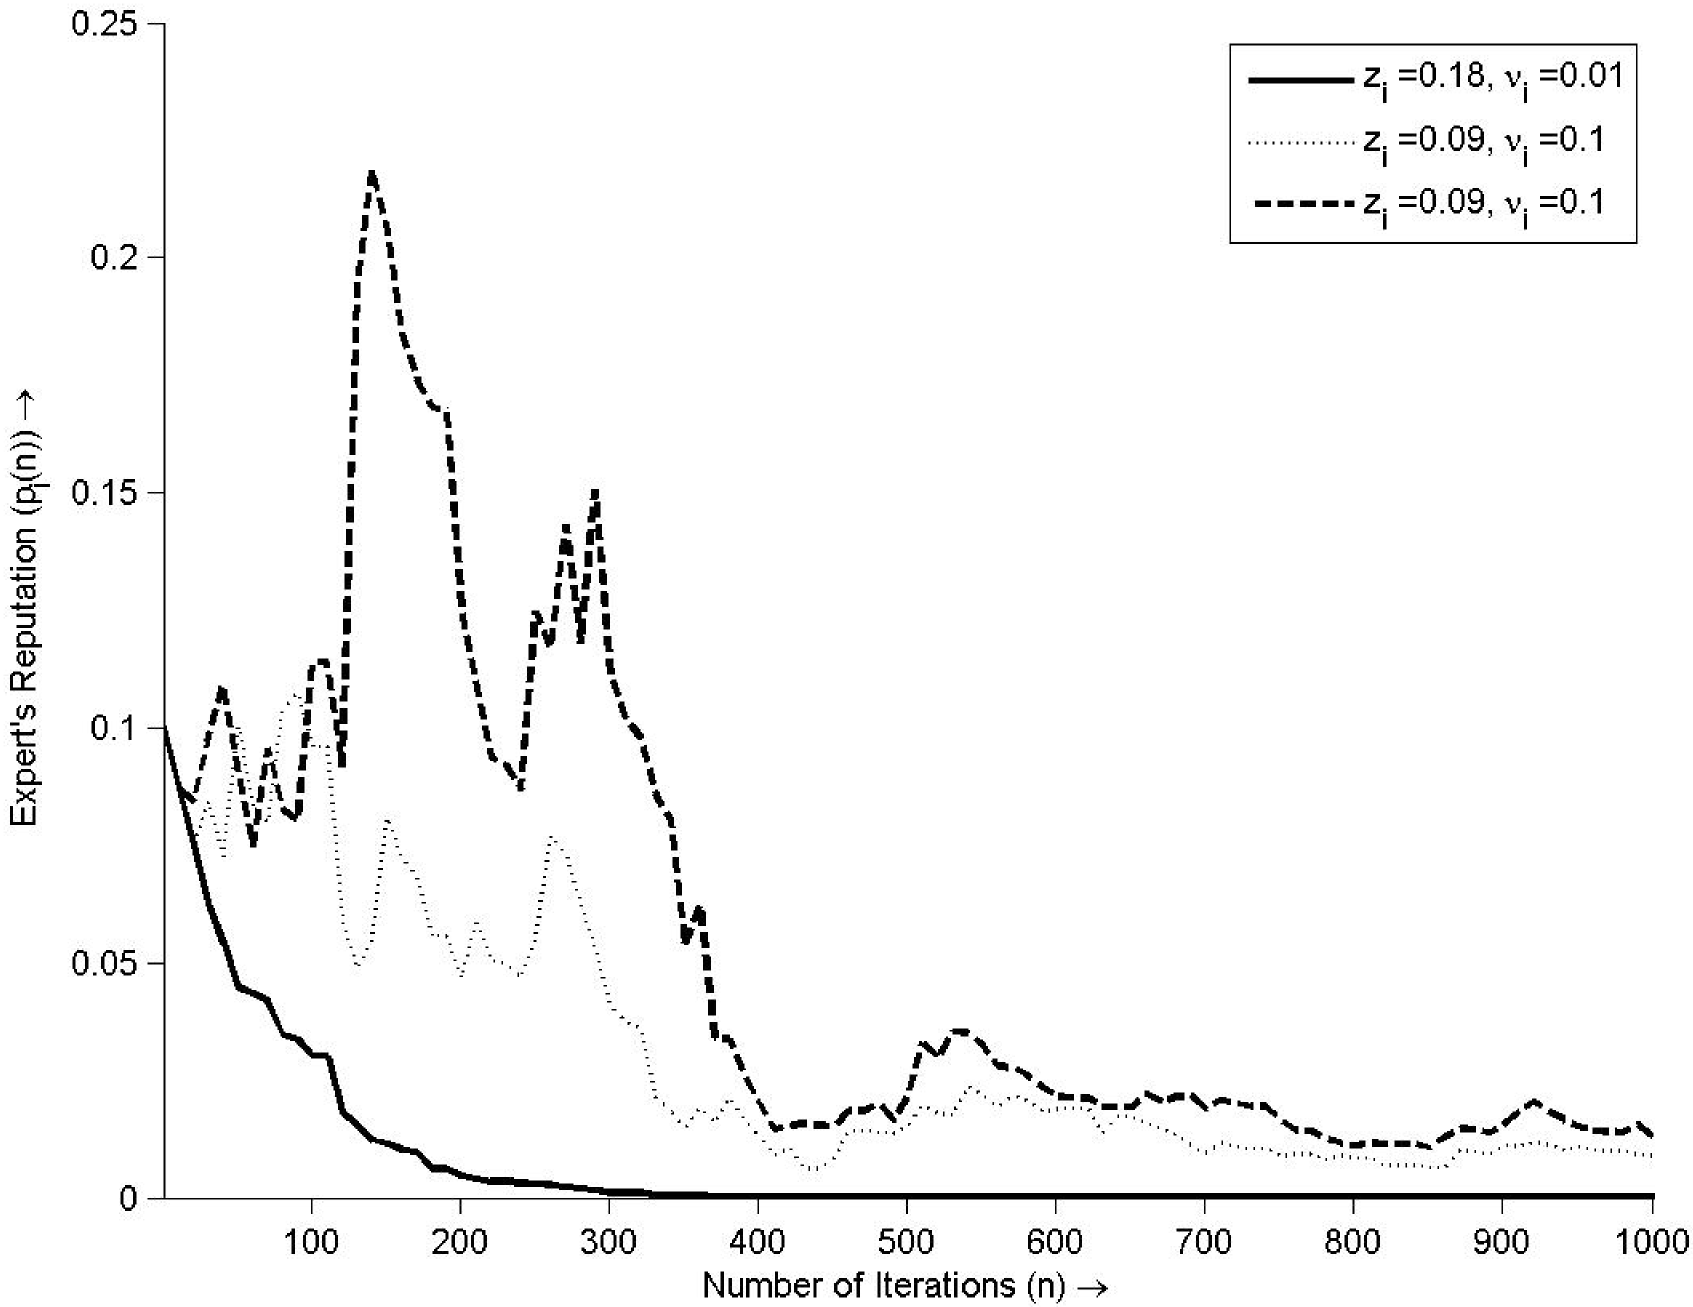

0.25
0.2
0.15
0.1
0.05
0
100
200
300
400
500
600
700
800
900
1000
Expert's Reputation ($\rho_i(n)$) →
Number of Iterations (n) →
z_i =0.18, ν_i =0.01
z_i =0.09, ν_i =0.1
z_i =0.09, ν_i =0.1

Steady State Behavior of Best 3 Experts' Reputation

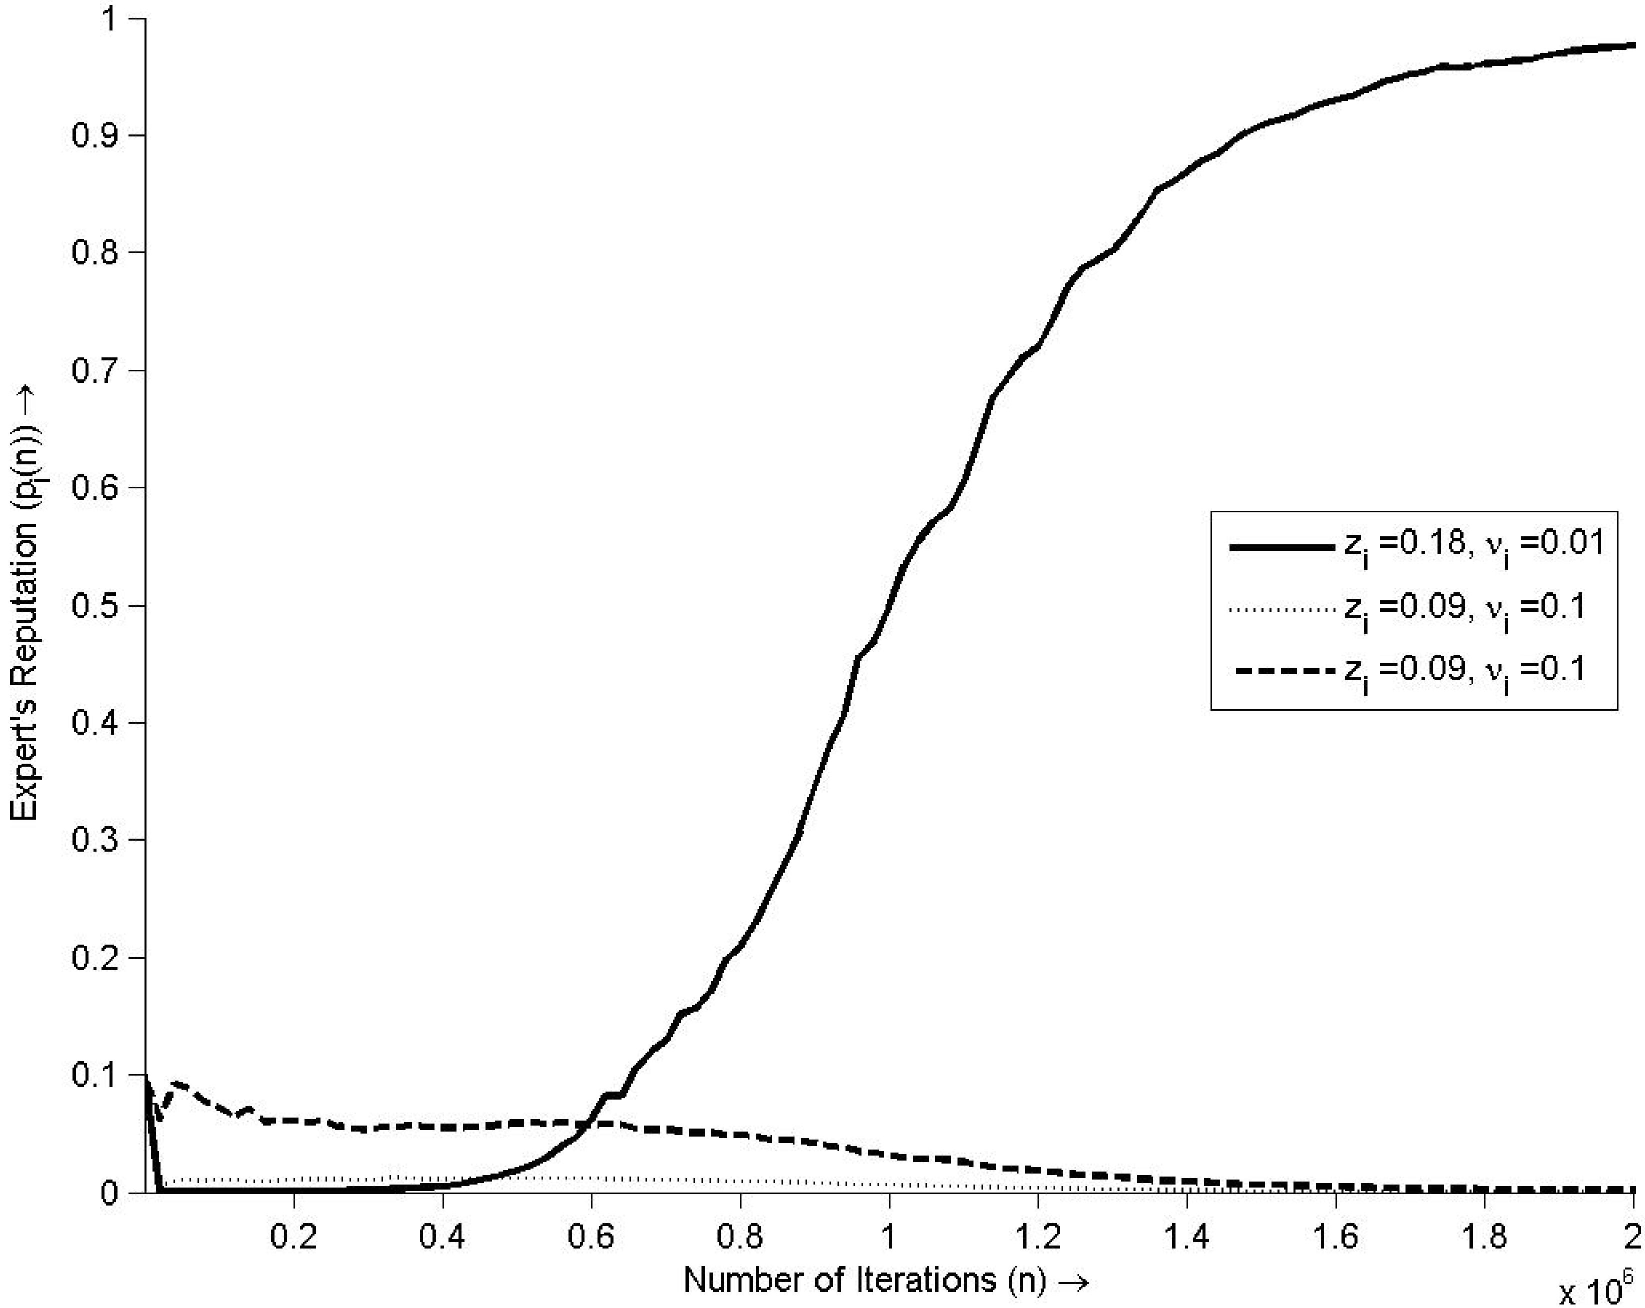

1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Expert's Reputation (p_i(n)) →
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
Number of Iterations (n) →
x 10^6
z_i =0.18, ν_i =0.01
z_i =0.09, ν_i =0.1
z_i =0.09, ν_i =0.1