A. Supplementary Material

A.1. Pseudocode for the generic CB algorithm and the DCB algorithm

Algorithm 2 Confidence Ball

Initialization: Set $A_0 = I$ and $b_0 = 0$.
for $t = 0, \dots \infty$ **do**
  Receive action set $\mathcal{D}_t$
  Construct the confidence ball C_t using A_t and b_t
  Choose action and receive reward:
    Find $(x_t, *) = \arg\max_{(x,\tilde{\theta}) \in \mathcal{D}_t \times C_t} x^\intercal \tilde{\theta}$
    Get reward r_t^i from context x_t^i
    Update $A_{t+1} = A_t + x_t x_t^\intercal$ and $b_{t+1} = b_t + r_t x_t$
end for

Algorithm 3 Distributed Confidence Ball

Input: Network V of agents, the function $\tau : t \to t - 4\log_2(|V|^{\frac{3}{2}} t)$.
Initialization: For each i, set $\tilde{A}_0^i = I_d$ and $\tilde{b}_0^i = \mathbf{0}$, and the buffers $\mathcal{A}_0^i = \emptyset$ and $\mathcal{B}_0^i = \emptyset$.
for $t = 0, \dots \infty$ **do**
  Draw a random permutation σ of $\{1, \dots, |V|\}$
  for each agent $i \in V$ **do**
    Receive action set $\mathcal{D}_t^i$ and construct the confidence ball C_t^i using $\tilde{A}_t^i$ and $\tilde{b}_t^i$
    Choose action and receive reward:
      Find $(x_{t+1}^i, *) = \arg\max_{(x,\tilde{\theta}) \in \mathcal{D}_t^i \times C_t^i} x^\intercal \tilde{\theta}$
      Get reward r_{t+1}^i from context x_{t+1}^i.
    Share and update information buffers:
      Set $\mathcal{A}_{t+1}^i = \left(\frac{1}{2}(\mathcal{A}_t^i + \mathcal{A}_t^{\sigma(i)})\right) \circ (x_{t+1}^i \left(x_{t+1}^i\right)^\intercal)$ and $\mathcal{B}_{t+1}^i = \left(\frac{1}{2}(\mathcal{B}_t^i + \mathcal{B}_t^{\sigma(i)})\right) \circ (r_{t+1}^i x_{t+1}^i)$
      if $|\mathcal{A}_{t+1}^i| > t - \tau(t)$ set $\tilde{A}_{t+1}^i = \tilde{A}_t^i + \mathcal{A}_{t+1}^i(1)$ and $\mathcal{A}_{t+1}^i = \mathcal{A}_{t+1}^i \setminus \mathcal{A}_{t+1}^i(1)$. Similary for $\mathcal{B}_{t+1}^i$.
  end for
end for

A.2. More on Communication Complexity

First, recall that if the agents want to communicate their information to each other at each round without a central server, then every agent would need to communicate their chosen action and reward to every other agent at each round, giving a communication cost of $O(d|V|^2)$ bits per-round. Under DCB each agent requires at most $O(log_2(|V|t)d^2|V|)$ bits to be communicated per round. Therefore, a significant communication cost reduction is gained when $log(|V|t)d \ll |V|$.

Recall also that using an epoch-based approach, as in (Szörényi et al., 2013), we reduce the per-round communication cost of the gossip-based approach to $O(d^2|V|)$. This makes the algorithm more efficient over any time horizon, requiring only that $d \ll |V|$, and the proofs of the regret performance are simple modifications of the proofs for DCB. In comparison with growing buffers this is only an issue after $O(\exp(|V|))$ number of rounds, and typically $|V|$ is large. This is why we choose to exhibit the growing-buffer approach in this current work.

Instead of relying on the combination of the diffusion and a delay to handle the potential doubling of data points under the randomised gossip protocol, we could attempt to keep track which observations have been shared with which agents, and thus simply stop the doubling from occurring. However, the per-round communication complexity of this is at least quadratic in $|V|$, whereas our approach is linear. The reason for the former is that in order to be efficient, any agent j, when sending information to an agent i, needs to know for each k which are the latest observations gathered by agent k that agent i already knows about. The communication cost of this is of order $|V|$. Since every agent shares information with somebody in each round, this gives per round communication complexity of order $|V|^2$ in the network.

A simple, alternative approach to the gossip protocol is a *Round-Robin* (RR) protocol, in which each agent passes the information it has gathered in previous rounds to the next agent in a pre-defined permutation. Implementing a RR protocol leads to the agents performing a distributed version of the CB-*InstSharing* algorithm, but with a delay that is of size at least linear in $|V|$, rather than the logarithmic dependence on this quantity that a gossip protocol achieves. Indeed, at any time, each agent will be lacking $|V|(|V|-1)/2$ observations. Using this observation, a cumulative regret bound can be achieved using Proposition 2 which arrives at the same asymptotic dependence on $|V|$ as our gossip protocol, but with an additive constant that is worse by a multiplicative factor of $|V|$. This makes a difference to the performance of the network when $|V|$ is very large. Moreover, RR protocols do not offer the simple generalisability and robustness that gossip protocols offer.

Note that the pruning protocol for DCCB only requires sharing the estimated θ-vectors between agents, and adds at most $O(d|V|)$ to the communication cost of the algorithm. Hence the per-round communication cost of DCCB remains $O(log_2(|V|t)d^2|V|)$.

Algorithm	Regret Bound	Per-Round Communication Complexity
CB-*NoSharing*	$O(\|V\|\sqrt{t})$	0
CB-*InstSharing*	$O(\sqrt{\|V\|t})$	$O(d\|V\|^2)$
DCB	$O(\sqrt{\|V\|t})$	$O(log_2(\|V\|t)d^2\|V\|)$
DCCB	$O(\sqrt{\|U^k\|t})$	$O(log_2(\|V\|t)d^2\|V\|)$

Figure 2. This table gives a summary of theoretical results for the multi-agent linear bandit problem. Note that CB with no sharing cannot benefit from the fact that all the agents are solving the same bandit problem, while CB with instant sharing has a large communication-cost dependency on the size of the network. DCB succesfully achieves near-optimal regret performance, while simultaneously reducing communication complexity by an order of magnitude in the size of the network. Moreover, DCCB generalises this regret performance at not extra cost in the order of the communication complexity.

A.3. Proofs of Intermediary Results for DCB

Proof of Proposition 1. This follows the proof of Theorem 2 in (Abbasi-Yadkori et al., 2011), substituting appropriately weighted quantities.

For ease of presentation, we define the shorthand

$$\tilde{X} := (\sqrt{w_1}y_1, \ldots, \sqrt{w_n}y_n) \text{ and } \tilde{\eta} = (\sqrt{w_1}\eta_1, \ldots, \sqrt{w_n}\eta_n)^\intercal,$$

where the y_i are vectors with norm less than 1, the η_i are R-subgaussian, zero mean, random variables, and the w_i are positive real numbers. Then, given samples $(\sqrt{w_1}y_1, \sqrt{w_1}(\theta y_1 + \eta_1)), \ldots, (\sqrt{w_n}y_n, \sqrt{w_n}(\theta y_n + \eta_n))$, the maximum likelihood estimate of θ is

$$\begin{aligned}\tilde{\theta} :&= (\tilde{X}\tilde{X}^\intercal + I)^{-1}\tilde{X}(\tilde{X}^\intercal\theta + \tilde{\eta}) \\ &= (\tilde{X}\tilde{X}^\intercal + I)^{-1}\tilde{X}\tilde{\eta} + (\tilde{X}\tilde{X}^\intercal + I)^{-1}(\tilde{X}\tilde{X}^\intercal + I)\theta - (\tilde{X}\tilde{X}^\intercal + I)^{-1}\theta \\ &= (\tilde{X}\tilde{X}^\intercal + I)^{-1}\tilde{X}\tilde{\eta} + \theta - (\tilde{X}\tilde{X}^\intercal + I)^{-1}\theta\end{aligned}$$

So by Cauchy-Schwarz, we have, for any vector x,

$$\begin{aligned}x^\intercal(\tilde{\theta} - \theta) &= \langle x, \tilde{X}\tilde{\eta}\rangle_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}} - \langle x, \theta\rangle_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}} && (14)\\ &\leq \|x\|_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}}\left(\|\tilde{X}\tilde{\eta}\|_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}} + \|\theta\|_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}}\right) && (15)\end{aligned}$$

Now from Theorem 1 of (Abbasi-Yadkori et al., 2011), we know that with probability $1-\delta$

$$\|\tilde{X}\tilde{\eta}\|^2_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}} \leq W^2R^2 2\log\sqrt{\frac{\det(\tilde{X}\tilde{X}^\intercal + I)}{\delta^2}}.$$

where $W = \max_{i=1,\ldots,n} w_i$. So, setting $x = (\tilde{X}\tilde{X}^\intercal + I)^{-1}(\tilde{\theta} - \theta)$, we obtain that with probability $1-\delta$

$$\|\tilde{\theta} - \theta\|_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}} \leq WR\left(2\log\sqrt{\frac{\det(\tilde{X}\tilde{X}^\intercal + I)}{\delta^2}}\right)^{\frac{1}{2}} + \|\theta\|_2$$

since [3]

$$\|x\|_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}}\|\theta\|_{(\tilde{X}\tilde{X}^\intercal+I)^{-1}} \le \|x\|_2\lambda_{\min}^{-1}(\tilde{X}\tilde{X}^\intercal+I)\|\theta\|_2\lambda_{\min}^{-1}(\tilde{X}\tilde{X}^\intercal+I)$$
$$\le \|x\|_2\|\theta\|_2.$$

Conditioned on the values of the weights, the statement of Proposition 1 now follows by substituting appropriate quantities above, and taking the probability over the distribution of the subGaussian random rewards. However, since this statement holds uniformly for any values of the weights, it holds also when the probability is taken over the distribution of the weights. □

Proof of Lemma 3. Recall that $\tilde{A}_t^i$ is constructed from the contexts chosen from the first $\tau(t)$ rounds, across all the agents. Let i' and t' be arbitrary indices in V and $\{1,\ldots,\tau(t)\}$, respectively.

(i) We have

$$\begin{aligned}\det\left(\tilde{A}_t^i\right) &= \det\left(\tilde{A}_t^i - \left(w_{i,t}^{i',t'}-1\right)x_{t'}^{i'}\left(x_{t'}^{i'}\right)^\intercal + \left(w_{i,t}^{i',t'}-1\right)x_{t'}^{i'}\left(x_{t'}^{i'}\right)^\intercal\right)\\ &= \det\left(\tilde{A}_t^i - \left(w_{i,t}^{i',t'}-1\right)x_{t'}^{i'}\left(x_{t'}^{i'}\right)^\intercal\right)\\ &\quad\cdot\left(1+\left(w_{i,t}^{i',t'}-1\right)\|x_{t'}^{i'}\|_{\left(\tilde{A}_t^i-\left(w_{i,t}^{i',t'}-1\right)x_{t'}^{i'}\left(x_{t'}^{i'}\right)^\intercal\right)^{-1}}\right)\end{aligned}$$

The second equality follows using the identity $\det(I + cB^{1/2}xx^\intercal B^{1/2}) = (1 + c\|x\|_B)$, for any matrix B, vector x, and scalar c. Now, we repeat this process for all $i' \in V$ and $t' \in \{1,\ldots,\tau(t)\}$ as follows. Let $(t_1,i_1),\ldots,(t_{|V|\tau(t)},i_{|V|\tau(t)})$ be an arbitrary enumeration of $V\times\{1,\ldots,\tau(t)\}$, let $B_0 = \tilde{A}_t^i$, and $B_s = B_{s-1} - (w_{i,t}^{i_s,t_s}-1)x_{t_s}^{i_s}\left(x_{t_s}^{i_s}\right)^\intercal$ for $s = 1,\ldots,|V|\tau(t)$. Then $B_{|V|\tau(t)} = A_{\tau(t)}$, and by the calculation above we have

$$\begin{aligned}\det\left(\tilde{A}_t^i\right) &= \det\left(A_{\tau(t)}\right)\prod_{s=1}^{|V|\tau(t)}\left(1+\left(w_{i,t}^{i_s,t_s}-1\right)\|x_{t_s}^{i_s}\|_{(B_s)^{-1}}\right)\\ &\le \det\left(A_{\tau(t)}\right)\exp\left(\sum_{s=1}^{|V|\tau(t)}\left(w_{i,t}^{i_s,t_s}-1\right)\|x_{t_s}^{i_s}\|_{(B_s)^{-1}}\right)\\ &\le \exp\left(\sum_{t'=1}^{\tau(t)}\sum_{i'=1}^{|V|}\left|w_{i,t}^{i',t'}-1\right|\right)\det\left(A_{\tau(t)}\right)\end{aligned}$$

(ii) Note that for vectors x, y and a matrix B, by the Sherman-Morrison Lemma, and Cauchy-Schwarz inequality we have that:

$$\begin{aligned}x^\intercal(B+yy^\intercal)^{-1}x = x^\intercal B^{-1}x - \frac{x^\intercal B^{-1}yy^\intercal B^{-1}x}{1+y^\intercal B^{-1}y} &\ge x^\intercal B^{-1}x - \frac{x^\intercal B^{-1}xy^\intercal B^{-1}y}{1+y^\intercal B^{-1}y}\\ &= x^\intercal B^{-1}x(1+y^\intercal B^{-1}y)^{-1}\end{aligned} \tag{16}$$

Taking

$$B = \left(\tilde{A}_t^i - \left(w_{i,t}^{i',t'}-1\right)x_{t'}^{i'}\left(x_{t'}^{i'}\right)^\intercal\right) \text{ and } y = \sqrt{w_{i,t}^{i',t'}-1}\,x_{t'}^{i'},$$

and using that $y^\intercal B^{-1}y \le \lambda_{min}(B)^{-1}y^\intercal y$, by construction, we have that, for any $t' \in \{1,\ldots,\tau(t)\}$ and $i' \in V$,

$$x^\intercal\left(\tilde{A}_t^i\right)^{-1}x \ge x^\intercal\left(\tilde{A}_t^i - \left(w_{i,t}^{i',t'}-1\right)x_{t'}^{i'}\left(x_{t'}^{i'}\right)^\intercal\right)^{-1}x(1+|w_{i,t}^{i',t'}-1|)^{-1}.$$

Performing this for each $i' \in V$ and $t' \in \{1,\ldots,\tau(t)\}$, taking the exponential of the logarithm and using that $\log(1+a) \le a$ like in the first part finishes the proof.

□

[3] $\lambda_{\min}(\,\cdot\,)$ denotes the smallest eigenvalue of its argument.

A.4. Proof of Theorem 6

Throughout the proof let i denote the index of some arbitrary but fixed agent, and k the index of its cluster.

Step 1: Show the true clustering is obtained in finite time. First we prove that with probability $1-\delta$, the number of times agents in different clusters share information is bounded. Consider the statements

$$\forall i, i' \in V,\ \forall t,\ \left(\|\hat{\theta}^i_{local,t} - \hat{\theta}^{i'}_{local,t}\| > c^{thresh}_{\lambda}(t) \right) \implies i' \notin U^k \tag{17}$$

and,

$$\forall t \geq C(\gamma, \lambda, \delta) = c^{thresh^{-1}}_{\lambda}\left(\frac{\gamma}{2}\right),\ i' \notin U^k,\ \|\hat{\theta}^i_{local,t} - \hat{\theta}^{i'}_{local,t}\| > c^{thresh}_{\lambda}(t). \tag{18}$$

where c^{thresh}_{λ} and A_{λ} are as defined in the main paper. Lemma 4 from (Gentile et al., 2014) proves that these two statements hold under the assumptions of the theorem with probability $1-\delta/2$.

Let i be an agent in cluster U^k. Suppose that (17) and (18) hold. Then we know that at time $t = \lceil C(\gamma, \lambda, \delta) \rceil$, $U^k \subset V^i_t$. Moreover, since the sharing protocol chooses an agent uniformly at random from V^i_t independently from the history before time t, it follows that the time until $V^i_t = U^k$ can be upper bounded by a constant $C = C(|V|, \delta)$ with probability $1-\delta/2$. So it follows that there exists a constant $C = C(|V|, \gamma, \lambda, \delta)$ such that the event

$$E := \{\text{(17) and (18) hold, and } (t \geq C(|V|, \gamma, \lambda, \delta) \implies V^i_t = U^k)\}$$

holds with probability $1-\delta$.

Step 2: Consider the properties of the weights after clustering. On the event E, we know that each cluster will be performing the algorithm DCB within its own cluster for all $t > C(\gamma, |V|)$. Therefore, we would like to directly apply the analysis from the proof of Theorem 1 from this point. In order to do this we need to show that the weights, $w^{i',t'}_{i,t}$, have the same properties after time $C = C(\gamma, |V|, \lambda, \delta)$ that are required for the proof of Theorem 1.

Lemma 7. *Suppose that agent i is in cluster U^k. Then, on the event E,*

(i) for all $t > C(|V|, \gamma, \lambda, \delta)$ and $i' \in V \setminus U^k$, $w^{i',t'}_{i,t} = 0$;

(ii) for all $t' \geq C(|V|, \gamma, \lambda, \delta)$ and $i' \in U^k$, $\sum_{i \in U^k} w^{i',t'}_{i,C(|V|,\gamma)} = |U^k|$;

(iii) for all $t \geq t' \geq C(|V|, \gamma, \lambda, \delta)$ and $i' \in U^k$, the weights $w^{i',t'}_{i,t}$, $i \in U^k$, are i.d..

Proof. See Appendix A.5. ☐

We must deal also with what happens to the information gathered before the cluster has completely discovered itself. To this end, note that we can write, supposing that $\tau(t) \geq C(|V|, \gamma, \lambda, \delta)$,

$$\tilde{A}^i_t := \sum_{i' \in U^k} \frac{w^{i',C}_{i,t}}{|U^k|} \tilde{A}^{i'}_C + \sum_{t'=C+1}^{\tau(t)} \sum_{i' \in U^k} w^{i',t'}_{i,t} x^{i'}_{t'} \left(x^{i'}_{t'}\right)^{\intercal}. \tag{19}$$

Armed with this observation we show that the fact that sharing within the appropriate cluster only begins properly after time $C = C(|V|, \gamma, \lambda, \delta)$ the influence of the bias is unchanged:

Lemma 8 (Bound on the influence of general weights)**.** *On the event E, for all $i \in V$ and t such that $T(t) \geq C(|V|, \gamma, \lambda, \delta)$,*

(i) $\det\left(\tilde{A}^i_t\right) \leq \exp\left(\sum_{t'=C}^{\tau(t)} \sum_{i' \in U^k} \left| w^{i',t'}_{i,t} - 1 \right|\right) \det\left(A^k_{\tau(t)}\right)$,

(ii) and $\|x^i_t\|^2_{\left(\tilde{A}^i_t\right)^{-1}} \leq \exp\left(\sum_{t'=C}^{\tau(t)} \sum_{i' \in U^k} \left| w^{i',t'}_{i,t} - 1 \right|\right) \|x^i_t\|^2_{\left(A^k_{\tau(t)}\right)^{-1}}$.

Proof. See Appendix A.5. □

The final property of the weights required to prove Theorem 1 is that their variance is diminishing geometrically with each iteration. For the analysis of DCB this is provided by Lemma 4 of (Szörényi et al., 2013), and, using Lemma 7, we can prove the same result for the weights after time $C = C(|V|, \gamma, \lambda, \delta)$:

Lemma 9. *Suppose that agent i is in cluster U^k. Then, on the event E, for all $t \geq C = C(|V|, \gamma, \lambda, \delta)$ and $t' < t$, we have*

$$\mathbb{E}\left((w_{i,t}^{j,t'} - 1)^2\right) \leq \frac{|U^k|}{2^{t-\max\{t',C\}}}.$$

Proof. Given the properties proved in Lemma 7, the proof is identical to the proof of Lemma 4 of (Szörényi et al., 2013). □

Step 3: Apply the results from the analysis of DCB. We can now apply the same argument as in Theorem 1 to bound the regret after time $C = C(\gamma, |V|, \lambda, \delta)$. The regret before this time we simply upper bound by $|U^k|C(|V|, \gamma, \lambda, \delta)\|\theta\|$. We include the modified sections bellow as needed.

Using Lemma 9, we can control the random exponential constant in Lemma 8, and the upper bound $W(T)$:

Lemma 10 (Bound in the influence of weights under our sharing protocol)**.** *Assume that $t \geq C(\gamma, |V|, \lambda\delta)$. Then on the event E, for some constants $0 < \delta_{t'} < 1$, with probability $1 - \sum_{t'=1}^{\tau(t)} \delta_{t'}$*

$$\sum_{t'=C}^{\tau(t)} \sum_{i' \in U^k} \left| w_{i,t}^{i',t'} - 1 \right| \leq |U^k|^{\frac{3}{2}} \sum_{t'=C}^{\tau(t)} \sqrt{\frac{2^{-(t-\max\{t',C\})}}{\delta_{t'}}},$$

$$\text{and } W(\tau(t)) \leq 1 + \max_{C \leq t' \leq \tau(t)} \left\{ |U^k|^{\frac{3}{2}} \sqrt{\frac{2^{-(t-\max\{t',C\})}}{\delta_{t'}}} \right\}.$$

In particular, for any $1 > \delta > 0$, choosing $\delta_{t'} = \delta 2^{-(t-\max\{t',C\})/2}$, and $\tau(t) = t - c_1 \log_2 c_2 t$ we conclude that with probability $1 - (c_2 t)^{-c_1/2}\delta/(1 - 2^{-1/2})$, for any $t > C + c_1 \log_2(c_2 C)$,

$$\sum_{i' \in U^k} \sum_{t'=C}^{\tau(t)} \left| w_{i,t}^{i',t'} - 1 \right| \leq \frac{|U^k|^{\frac{3}{2}} (c_2 t)^{-\frac{c_1}{4}}}{(1 - 2^{-\frac{1}{4}})\sqrt{\delta}}, \text{ and } W(\tau(t)) \leq 1 + \frac{|U^k|^{\frac{3}{2}} (c_2 t)^{-\frac{c_1}{4}}}{\sqrt{\delta}}. \tag{20}$$

Thus lemmas 8 and 10 give us control over the bias introduced by the imperfect information sharing. Applying lemmas 8 and 10, we find that with probability $1 - (c_2 t)^{-c_1/2}\delta/(1 - 2^{-1/2})$:

$$\rho_t^i \leq 2 \exp\left(\frac{|U^k|^{\frac{3}{2}}}{(1 - 2^{-\frac{1}{4}}) c_2^{\frac{c_1}{4}} t^{\frac{c_1}{4}} \sqrt{\delta}} \right) \|x_t^i\|_{\left(A_{\tau(t)}^i\right)^{-1}} \tag{21}$$

$$\cdot \left[\left(1 + \frac{|U^k|^{\frac{3}{2}}}{(1 - 2^{-\frac{1}{4}}) c_2^{\frac{c_1}{4}} t^{\frac{c_1}{4}} \sqrt{\delta}} \right) \left[R \sqrt{2 \log \left(\exp\left(\frac{|U^k|^{\frac{3}{2}}}{(1 - 2^{-\frac{1}{4}}) c_2^{\frac{c_1}{4}} t^{\frac{c_1}{4}} \sqrt{\delta}} \right) \frac{\det\left(A_{\tau(t)}\right)^{\frac{1}{2}}}{\delta} \right)} + \|\theta\| \right] \right].$$

Step 4: Choose constants and sum the simple regret. Choosing again $c_1 = 4$, $c_2 = |V|^{\frac{3}{2}}$, and setting $N_\delta = \frac{1}{(1-2^{-\frac{1}{4}})\sqrt{\delta}}$, we have on the event E, for all $t \geq \max\{N_\delta, C + 4\log_2(|V|^{\frac{3}{2}} C)\}$, with probability $1 - (|V|t)^{-2}\delta/(1 - 2^{-1/2})$

$$\rho_t^i \leq 4e \|x_t^i\|_{\left(A_{t-1}^k + \sum_{i'=1}^{i-1} x_t^{i'} \left(x_t^{i'}\right)^\intercal\right)^{-1}} \left(\beta(t) + R\sqrt{2} \right),$$

where $\beta(\cdot)$ is as defined in the theorem statement. Now applying Cauchy-Schwarz, and Lemma 11 from (Abbasi-Yadkori et al., 2011) yields that on the event E, with probability $1 - \left(1 + \sum_{t=1}^{\infty} (|V|t)^{-2}/(1 - 2^{-1/2})\right)\delta \geq 1 - 3\delta$,

$$\mathcal{R}_t \leq \left(\max\{N_\delta, C + 4\log_2(|V|^{\frac{3}{2}} C)\} + 2\left(4|V| d \log\left(|V|t\right)\right)^3 \right) \|\theta\|_2$$

$$+ 4e\left(\beta(t) + R\sqrt{2}\right)\sqrt{|U^k|t\left(2\log\left(\det\left(A_t^k\right)\right)\right)}.$$

Replacing δ with $\delta/6$, and combining this result with Step 1 finishes the proof.

A.5. Proofs of Intermediary Results for DCCB

Proof of Lemma 7. Recall that whenever the pruning procedure cuts an edge, both agents reset their buffers to their local information, scaled by the size of their current neighbour sets. (It does not make a difference practically whether or not they scale their buffers, as this effect is washed out in the computation of the confidence bounds and the local estimates. However, it is convenient to assume that they do so for the analysis.) Furthermore, according to the pruning procedure, no agent will share information with another agent that does not have the same local neighbour set.

On the event E, there is a time for each agent, i, before time $C = C(\gamma, |V|, \lambda\delta)$ when the agent resets its information to their local information, and their local neighbour set becomes their local cluster, i.e. $V_t^i = U^k$. After this time, this agent will only share information with other agents that have also set their local neighbour set to their local cluster. This proves the statement of part (i).

Furthermore, since on event E, after agent i has identified its local neighbour set, i.e. when $V_t^i = U^k$, the agent only shares with members of U^k, the statements of parts (ii) and (iii) hold by construction of the sharing protocol. □

Proof of Lemma 8. The result follows the proof of Lemma 3. For the the iterations until time $C = C(\gamma, |V|, \lambda\delta)$ is reached, we apply the argument there. For the final step we require two further inequalities.

First, to finish the proof of part (i) we note that,

$$\det\left((A_T^k - A_C^k) + \sum_{i' \in U^k} \frac{w_{i,t}^{i',C(\gamma,|V|)}}{|U^k|}\tilde{A}_C^{i'}\right) = \det\left(A_T^k + \sum_{i' \in U^k} \frac{w_{i,t}^{i',C} - 1}{|U^k|}\tilde{A}_C^{i'}\right)$$

$$= \det\left(A_T^k\right)\det\left(I + \sum_{i' \in U^k} \frac{w_{i,t}^{i',C} - 1}{|U^k|} {A_T^k}^{-\frac{1}{2}}\tilde{A}_C^{i'}{A_T^k}^{-\frac{1}{2}}\right)$$

$$\leq \det\left(A_T^k\right)\det\left(I + \left[\sum_{i' \in U^k}\left|w_{i,t}^{i',C} - 1\right|\right]{A_T^k}^{-\frac{1}{2}}\sum_{i' \in U^k}\frac{\tilde{A}_C^{i'}}{|U^k|}{A_T^k}^{-\frac{1}{2}}\right) \leq \det\left(A_T^k\right)\left(1 + \sum_{i' \in U^k}\left|w_{i,t}^{i',C} - 1\right|\right).$$

For the first equality we have used that $|U^k|A_C^k = \sum_{i' \in U^k}\tilde{A}_C^{i'}$; for the first inequality we have used a property of positive definite matrices; for the second inequality we have used that 1 upper bounds the eigenvalues of ${A_T^k}^{-1/2}A_C^k{A_T^k}^{-1/2}$.

Second, to finish the proof of part (ii), we note that, for any vector x,

$$x^\intercal\left(A_{\tau(t)}^k + \sum_{i' \in U^k}\frac{w_{i,t}^{i',C} - 1}{|U^k|}\tilde{A}_C^{i'}\right)^{-1} x$$

$$= \left({A_{\tau(t)}^k}^{-\frac{1}{2}}x\right)^\intercal\left(I + \sum_{i' \in U^k}\frac{w_{i,t}^{i',C} - 1}{|U^k|}{A_{\tau(t)}^k}^{-\frac{1}{2}}\tilde{A}_C^{i'}{A_{\tau(t)}^k}^{-\frac{1}{2}}\right)^{-1}\left({A_{\tau(t)}^k}^{-\frac{1}{2}}x\right)$$

$$\geq \left({A_{\tau(t)}^k}^{-\frac{1}{2}}x\right)^\intercal\left(I + \sum_{i' \in U^k}\frac{\left|w_{i,t}^{i',C} - 1\right|}{|U^k|}{A_{\tau(t)}^k}^{-\frac{1}{2}}\tilde{A}_C^{i'}{A_{\tau(t)}^k}^{-\frac{1}{2}}\right)^{-1}\left({A_{\tau(t)}^k}^{-\frac{1}{2}}x\right)$$

$$\geq \left(1 + \sum_{i' \in U^k}\left|w_{i,t}^{i',C} - 1\right|\right)^{-1} x^\intercal {A_{\tau(t)}^k}^{-1}x.$$

The first inequality here follows from a property of positive definite matrices, and the other steps follow similarly to those in the inequality that finished part (i) of the proof. □