
Train 'n Trade: Foundations of Parameter Markets

Tzu-Heng Huang, Harit Vishwakarma, Frederic Sala

Department of Computer Science

University of Wisconsin-Madison

{thuang273, hvishwakarma}@wisc.edu, fredsa@cs.wisc.edu

Abstract

Organizations typically train large models individually. This is costly and time-consuming, particularly for large-scale foundation models. Such vertical production is known to be suboptimal. Inspired by this economic insight, we ask whether it is possible to leverage others' expertise by trading the constituent parts in models, i.e., sets of weights, as if they were market commodities. While recent advances in aligning and interpolating models suggest that doing so may be possible, a number of fundamental questions must be answered to create viable parameter markets. In this work, we address these basic questions, propose a framework containing the infrastructure necessary for market operations to take place, study strategies for exchanging parameters, and offer means for agents to monetize parameters. Excitingly, compared to agents who train siloed models from scratch, we show that it is possible to mutually gain by using the market, even in competitive settings. This suggests that the notion of parameter markets may be a useful paradigm for improving large-scale model training in the future.

1 Introduction

Costs to build powerful state-of-the-art machine learning models, such as foundation models (e.g., GPT-3 [1], T5 [2], PaLM [3], BLOOM-176B [4] and OPT-175B [5]), have increased enormously. These costs can easily reach millions of dollars and even exceed that amount. For example, training the 11 billion-parameter T5 model is estimated to take around 1.3 million dollars for a single training run [6]. Unfortunately, few organizations and fewer individuals are sufficiently well-capitalized to afford such training costs.

One approach to reduce expense is to broadly distribute training workloads, such as in decentralized training [7, 8, 9, 10]. However, this is limiting; even in the decentralized setting, participants must first agree to train a shared model and at least minimally coordinate the training process. For this reason, such techniques cannot be applied when organizations develop different models for different purposes on different timelines. In these scenarios—the most common solution in large-scale model development—models are trained individually regardless of high cost.

A natural question is whether such *vertical production* can be broken down into parts that can be more easily built, re-used, and exchanged. Taking inspiration from other areas of manufacturing, we observe that most products are not built in vertical silos, but from components that are traded in markets. Economic agents, even when competing against each other, buy, sell, and trade such components to leverage the expertise of other agents so that production costs can be lowered.

This leads us to ask whether subsets of trained weights can be thought of as constituent parts to be bought and sold on *parameter markets*. Such markets may provide mutual benefits for both buyers and sellers. Buyers are able to purchase well-trained parameter sets directly as commodities to leverage the training expertise of others and then use them to improve model performance. Sellers (i.e., owners of partially or fully-trained models) are able to monetize parameters as a second profit center, in addition to the downstream activity enabled by using their models.

Challenges and Proposed Approach. How can we build and use such markets? To answer this, we must first overcome several obstacles. An immediate challenge is the notion of *alignment*: models trained in isolation may not have parameter sets that correspond in any natural way, especially if these models have differing purposes. Excitingly, recent work suggests that it is possible to align model components and then merge them via linear interpolation [11, 12, 13, 14, 15, 16, 17]. Similarly, it is known that training data can be potentially recovered from weights or gradients so privacy is an additional challenge [18, 19, 20]. We tackle the orthogonal question:

If model alignment and sufficient privacy can indeed be assured, how can a viable parameter marketplace be designed?

There are three fundamental challenges in doing so:

1. *How should agents (users in the market) decide to perform transactions?* Discovering and verifying useful model parameter sets on the market without prior knowledge is challenging for buyers. To address this issue, we introduce a trusted third-party organization, the *broker*. The broker enables a "try-before-purchase" mechanism for buyers to examine the quality of parameters. This is a common approach in the existing works on *data* marketplaces [21, 22] as it allows buyers to evaluate the quality of data before purchase. Doing so with parameters rather than data presents additional complications that must be resolved.
2. *What rewards may agents gain?* An important consideration when using such a framework is whether *agents can expect to see any gains*. Indeed, if the process of exchanging and validating parameter sets does not yield any improvements when compared to siloed training, there is no reason to participate in parameter markets. We provide theoretical and empirical analyses validating the improvements gained from participating in such markets.
3. *How to monetize model parameters in a competitive market?* In settings where parameters are bought and sold, rather than bartered, it can be challenging to price these assets. Both the seller and the buyer may not have a clear understanding of the other's valuation of the parameters, which makes it difficult for each to maximize their revenues in a trade. To address this issue, we apply a Bayesian-optimal pricing mechanism [23] to provide valuation of parameter sets and study Nash bargaining [24] to find market prices in negotiation.

Results and Contributions. We propose and formulate a viable marketplace to trade parameters for machine learning models. We validate it in both theoretical and empirical settings. Theoretically, in basic scenarios we show how agent training converges faster through purchasing parameters in the market. We offer bounds on the improvement gained via trading when training linear models. Empirically, we conduct experiments in a variety of practical scenarios to validate the framework's effectiveness. We demonstrate that compared to agents who stay outside the market and train models in isolation, participating in parameter trading, even trading subsets of the full set of model parameters, provides benefits to efficient model training and better model performance. For example, when training and trading parameters of ResNet20 on TinyImageNet, two agents improve their performance by gaining accuracy improvements of **+10.02%** and **+15.93%** versus separate training. We also demonstrate the success of price estimation to monetize parameters and negotiate prices.

2 Related Works

First, we describe two related concepts: *data* and *model*—as opposed to *parameter*—marketplaces. We then give background on model alignment techniques, which are used in our framework.

Data Marketplaces. Data is a key ingredient in machine learning pipelines. There is a rich vein of work proposing infrastructure to trade data as a commodity [21, 22, 25, 26, 27, 28, 29]. Such marketplaces have also been instantiated in industry, including in services such as Amazon Web Services (AWS) Data Exchange, Microsoft's Azure Data Marketplace, and Google's Cloud Data Catalog. Such research; however, cannot be directly applied to trading parameters. It is relatively easy to evaluate the valuation of a data trade using basic statistical measurements. In contrast, the value of parameters is challenging to measure, as it can only be determined after testing and depends on the model's performance.

Model Marketplaces. There have been several efforts to build markets to trade *entire model instances* trained by a centralized broker [30, 31]. Two major obstacles that need to be overcome are determining the value and pricing models, and safeguarding privacy. To address the former issue, a noise-injection mechanism has been suggested, which involves assessing the accuracy of an optimal model with random Gaussian noise to determine its worth and creating a price-error curve for selling it on the market [30]. The latter issue has been tackled by proposing a system that apply differential privacy while still maximizing revenue [31]. In contrast to trading entire model instances for downstream use, parameter markets are far more refined, enabling each user in the market to train their own models for their own purposes while gaining from others’ training runs.

Model Alignment. Models that are trained with different batch orders or initialization weights may not be properly aligned. Directly merging purchased parameters through interpolation may fail. The process of aligning parameters in model training is therefore critical. Recent studies have explored the geometric relationship between models [11, 12, 13, 14, 15] and proposed methods for aligning two sets of neural network parameters [16, 17]. We use such techniques as a building block in our proposed market infrastructure. Through proper model alignments, we expect agents are able to find and purchase desired parameter sets in the market.

3 A Framework for Parameter Markets

We provide a general description of the proposed marketplace and then discuss each component in depth.

3.1 General Marketplace Framework

Figure 1 depicts a two-agent version of the marketplace. Multiple agents training models for potentially different tasks seek to buy or sell sets of parameters. Buying well-trained parameter sets enables reaching a goal performance faster while selling such parameters produces profits.

A high-level description of the trading process follows. First, agents send their own parameters to the market ①. A third party (the broker) operates a "try-before-purchase" routine to align and merge parameters for buyers ②. The broker privately informs agents of the gains or losses resulting from a potential trade using validation data ③. Based on this information, a buyer values a seller’s parameters, and then makes a trading decision. If the buyer is willing to purchase, a quote request is sent out, and both sides generate and report their valuations to the broker ④, ⑤. The broker helps both parties negotiate the price until they reach an agreement ⑥. Afterwards, the broker ships parameters to the buyer and transfers the payment to the seller, completing the trade.

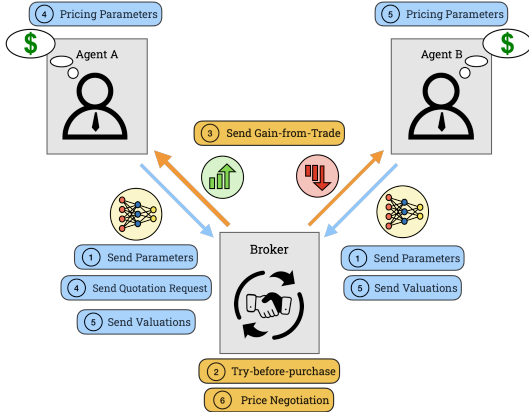


Figure 1: Overall workflow in a two-agent market. Blue and orange blocks represent actions taken by agents and the broker, respectively. In this example, agent *A* is informed of a potential gain through purchasing agent *B*’s parameters. Hence, agent *A* sends a quotation request to inquire about purchasing parameters. Then, broker helps both sides negotiate on the price of agent *B*’s parameters.

3.2 Market Setup

In the following sections, we discuss each foundational concept in parameter markets. We first fix some notation. For agent u , let $D_u = \{s_{u,i}\}_{i=1}^{n_u}$ be the samples drawn from a distribution \mathcal{D}^u supported on \mathcal{S} . At round t , let $\theta_u^t \in \mathbb{R}^d$ be trained parameters that agent u has access to, and let $\hat{\mathcal{L}}_u(\theta_u^t)$ be the empirical loss of agent u for the corresponding model measured on data D_u . The

empirical loss is defined with the agents' loss function ℓ in a standard fashion,

$$\hat{\mathcal{L}}_u(\theta_u^t) := \frac{1}{n_u} \sum_{i=1}^{n_u} \ell(\theta_u^t, s_{u,i}).$$

While our framework can handle general settings (i.e., any well-trained parameter sets can be traded), we focus on supervised learning for ease of exposition. We assume $\mathcal{S} = \mathcal{X} \times \mathcal{Y}$ where \mathcal{X} and \mathcal{Y} are the instance and label spaces respectively. Thus, $s_{u,i} = (x_{u,i}, y_{u,i})$, and also denote D_u as (X_u, Y_u) where X_u is the set of points $x_{u,i}$ and Y_u is the set of corresponding labels. We drop the superscript u when it is otherwise clear from the context.

Agents and Broker. For simplicity, suppose there are two agents A, B in the market. Agent A has n_a data points (X_a, Y_a) , and agent B has n_b data points (X_b, Y_b) . There is also a *trusted* third party, the broker Z , helping the two agents operate in the market impartially. The broker has access to a validation dataset (X_z, Y_z) of size n_z with noiseless samples from both distributions.

In the proposed trading framework, the two agents train their models locally first and then trade with each other. At the beginning of round t , agents perform a standard gradient descent update using their own datasets. That is, $\hat{\theta}_u^t := \theta_u^{t-1} - \eta \nabla \hat{\mathcal{L}}_u(\theta_u^{t-1})$, $u \in \{a, b\}$, where η indicates the step size.

Using Purchased Parameters. In the market, parameters can be bought or sold. To use acquired parameters, several model alignment techniques can be employed [16, 17]. For a potential trade, broker aligns seller's parameter sets to match a buyer's. Post-alignment, the parameters can be merged as a straightforward convex combination. For instance, if agent A is willing to buy parameters from agent B , and a trade occurs eventually, the post-alignment combination will have the form $\bar{\theta}_a^t := (1 - \alpha)\hat{\theta}_a^t + \alpha\hat{\theta}_b^t$, $\alpha \in (0, 1]$. Here α and (for agent B , β) are weights indicating the portion of trained parameters from the seller side in the merged parameters.

3.3 Try-Before-Purchase Mechanism

However, before making any trading decision, agents are unaware of the potential benefits they can achieve through purchasing. Our proposal is for a neutral broker to implement a "try-before-purchase" mechanism which assists agents. It does so by helping them align and merge parameters then pre-evaluate their quality by using the broker's dataset (X_z, Y_z) . In the try-before-purchase mechanism, the broker can pick the optimized weights α or β for buyer agents by minimizing the empirical loss. We write $\alpha := \arg \min_{\nu \in (0,1]} \hat{\mathcal{L}}_z((1 - \nu)\hat{\theta}_a^t + \nu\hat{\theta}_b^t)$ and $\beta := \arg \min_{\nu \in (0,1]} \hat{\mathcal{L}}_z((1 - \nu)\hat{\theta}_b^t + \nu\hat{\theta}_a^t)$.

Using the optimized purchased weights, the broker calculates and communicates to agents their gain or loss from the trade in a confidential manner. We denote the *gain-from-trade* for agent u by Δ_u^t , $u \in \{a, b\}$, which serves as prior knowledge to help agents make informed decisions about whether to make purchases or not. Generally, the notion of gain-from-trade is to compare relative improvement versus not trading. The trading benefits can be expressed in various ways, such as the difference between agent's loss before and after the trade. For example, it might take the form $\Delta_u^t = \hat{\mathcal{L}}_z(\hat{\theta}_u^t) - \hat{\mathcal{L}}_z(\bar{\theta}_u^t)$. Other ways to define the gain-from-trade include using the relative improvement ratio on the empirical loss or improvement ratio on the (estimated) parameter recovery error.

If the gain-from-trade Δ_u^t does not indicate any benefit for buyer agent u , the agent will not be willing to make a purchase. Consequently, no quote request will be sent out, and no trade will take place. In such a scenario, the parameters $\hat{\theta}_u^t$ will remain with agent u until the next round of gradient descent. The final parameters at the end of round t , denoted by θ_u^t , can be either $\bar{\theta}_u^t$ or $\hat{\theta}_u^t$. To indicate a trade, we use the indicator variable I_u^t :

$$\theta_u^t := (1 - I_u^t) \cdot \hat{\theta}_u^t + I_u^t \cdot \bar{\theta}_u^t \quad \text{where } I_u^t = \begin{cases} 1 & \text{if agent } u \text{ buys parameters,} \\ 0 & \text{if agent } u \text{ does not make a purchase.} \end{cases}$$

3.4 Valuation and Pricing Mechanism

Once the gain-from-trade has been determined and trading decisions have been made, quote requests with purchased weights are circulated throughout the market. Both buyers and sellers begin to assess various parameters in order to negotiate prices. Each agent has their own private valuation function, denoted by $v_u : \mathbb{R}^d \mapsto \mathbb{R}^+$, which quantifies their trading benefits produced by specific parameters $\theta \in \mathbb{R}^d$. The valuation function of each agent is private and is not exposed to others. For instance, agent A 's valuation function is denoted by $v_a(\hat{\theta}_b^t)$ for the value of agent B 's parameters for purchasing, and $v_a(\hat{\theta}_a^t)$ is the value of their own parameters for selling. In order to generate revenue from a trade, $v_a(\hat{\theta}_b^t)$ is the highest price that agent A is willing to bid for purchase, while $v_a(\hat{\theta}_a^t)$ is the lowest price that agent A is willing to ask for in order to sell.

Once valuations are completed, the broker—in their role as an impartial middleman—assists in bargaining the market price to maximize revenue for both sides. The negotiation process for prices continues until a mutually acceptable market price for parameters is reached. When a buyer offers to pay a lower price for certain parameters than the seller is willing to accept, the trade cannot be fulfilled. To analyze this negotiation process, we treat it as a Nash bargaining problem [24]. We assume that the broker is knowledgeable about the valuations of both parties (not the valuation functions themselves) and sets the price by maximizing the revenue of both agents impartially. An agent's revenue, defined as U_u , is derived from two sources: the profit earned from selling self-parameters and the profit gained from buying parameters. We employ the popular Cobb-Douglas function to maximize both agents' revenue [32]. We denote the market price for agent u 's parameters at round t as $P_u^t \in \mathbb{R}^+$, then formulate the problem accordingly. Set $U_a(P_a^t, P_b^t) := (P_a^t - v_a(\hat{\theta}_a^t)) + (v_a(\hat{\theta}_b^t) - P_b^t)$ and $U_b(P_b^t, P_a^t) := (P_b^t - v_b(\hat{\theta}_b^t)) + (v_b(\hat{\theta}_a^t) - P_a^t)$. Then we have the problem

$$\begin{aligned} & \operatorname{argmax}_{P_a^t, P_b^t} U_a(P_a^t, P_b^t) \times U_b(P_b^t, P_a^t) \\ \text{s.t.} \quad & P_a^t \in [v_a(\hat{\theta}_a^t), v_a(\hat{\theta}_b^t)], \quad P_b^t \in [v_b(\hat{\theta}_b^t), v_b(\hat{\theta}_a^t)]. \end{aligned}$$

By solving this problem, the broker can determine the difference in price, denoted by ΔP_{ab}^t , using

$$\Delta P_{ab}^t = P_a^t - P_b^t = \frac{1}{2} \left(v_b(\hat{\theta}_a^t) + v_a(\hat{\theta}_a^t) - v_a(\hat{\theta}_b^t) - v_b(\hat{\theta}_b^t) \right). \quad (1)$$

The steps for solving this problem are shown in the Appendix A. The resulting price difference ΔP_{ab}^t represents the amount of money that needs to be transferred between parties.

4 Instantiating the Market: Concrete Examples

Now we give a concrete instantiation of the market that we described in the previous section.

Valuations. There are various ways for agents to define valuation functions (i.e. based on agent's preference, training budget, or model performance). Here we assume that in the market, agents to purchase use gain-from-trade Δ_u^t , which can be seen as a notion of relative performance improvement, to assess the value of parameters so that $v_u(\hat{\theta}_{u'}^t) = \Delta_u^t$, where u' represents their seller agent.

However, assessing the value of self-parameters for agent u , who is a seller, is a difficult task as there is no clear information available from the broker regarding the quality of such parameters. The best approach for a seller to maximize profit is to set a price as close as possible to the buyer's valuation, which is the highest price that the buyer is willing to pay. To arrive at this *virtual valuation*, we use the Bayesian-optimal pricing mechanism described in Myerson's work [23]. This mechanism enables the seller to monetize self-parameters. Under this Bayesian mechanism, we assume that the seller is also aware that the buyer's valuation arises from gain-from-trade, and that their valuation function is derived from a known probability distribution. We discuss these common priors in Appendix G.

Suppose the buyer's valuation has a cumulative distribution function F_v . If the seller sets a price of P , the probability that the buyer is willing to purchase is $\mathbb{P}(P < v) = 1 - \mathbb{P}(v \leq P) = 1 - F_v(P)$. The expected revenue for the seller is $P \times (1 - F_v(P))$. Hence, the optimal price to ask for can be

Algorithm 1 Single Round of Parameter Trading

Input: $(X_a, Y_a), (X_b, Y_b), (X_z, Y_z), \theta^*, \theta_a^{t-1}, \theta_b^{t-1}$

Output: θ_b^t

$\hat{\theta}_u^t \leftarrow \theta_u^{t-1} - \eta \nabla \hat{\mathcal{L}}_u(\theta_u^{t-1}), u \in \{a, b\}$ ▷ agents' local training

$\bar{\theta}_a^t = (1 - \alpha)\hat{\theta}_a^t + \alpha\hat{\theta}_b^t, \quad \alpha = \arg \min_{\nu \in (0,1]} \hat{\mathcal{L}}_z((1 - \nu)\hat{\theta}_a^t + \nu\hat{\theta}_b^t)$ ▷ broker's try-before-purchase

$\bar{\theta}_b^t = (1 - \beta)\hat{\theta}_b^t + \beta\hat{\theta}_a^t, \quad \beta = \arg \min_{\nu \in (0,1]} \hat{\mathcal{L}}_z((1 - \nu)\hat{\theta}_b^t + \nu\hat{\theta}_a^t)$ ▷ broker's try-before-purchase

$\Delta_u^t = \frac{\|\hat{\theta}_u^t - \theta^*\|_2^2}{\|\bar{\theta}_u^t - \theta^*\|_2^2}, u \in \{a, b\}$ ▷ inform agents about gain-from-trade

if $\Delta_b^t > 1$ **then** ▷ agent B sends a quotation request with β

agent A, B provide valuations to broker ▷ agent A 's valuation is estimated by the bounds of Δ_b^t

if $v_b(\hat{\theta}_a^t) \geq v_a(\hat{\theta}_a^t)$ **then**

transfered payment for buying $\hat{\theta}_a^t$ after negotiation is set to $(v_b(\hat{\theta}_a^t) + v_a(\hat{\theta}_a^t))/2$

return $\bar{\theta}_b^t$ ▷ ship merged parameters

else

return $\hat{\theta}_b^t$ ▷ negotiation fails

else

return $\hat{\theta}_b^t$ ▷ agent B decides not to buy

found by maximizing the expected revenue, and it satisfies

$$P^* := \frac{1 - F_v(P^*)}{F'_v(P^*)}. \quad (2)$$

Linear Model Case Study. To further illustrate a concrete example of the seller's virtual valuation, we consider pricing parameters for training linear models. For simplicity, we assume that the broker knows the true parameters θ^* —though this is not necessary in practice, as these can be approximated using the broker's validation dataset. Both agents' data is sampled from the same distribution so that true parameters are identical. Additionally, we take the gain-from-trade revealed by broker to be the ratio of squared parameter estimation error. We write

$$\Delta_u^t := \frac{\|\hat{\theta}_u^t - \theta^*\|_2^2}{\|\bar{\theta}_u^t - \theta^*\|_2^2}. \quad (3)$$

If $\Delta_u^t > 1$, the agent's model is able to get closer to the true parameter (θ^*) by purchasing. We use agent A to illustrate the sale of their parameters $\hat{\theta}_a^t$ to agent B . First, we show how agent A determines bounds for the buyer's valuation $v_b(\hat{\theta}_a^t)$. Recall that once agent B expresses a willingness to purchase, a quote request with their purchased weight β will be communicated to agent A through the broker. Therefore, when agent A evaluates self-parameters, the broker provides three pieces of information: the buyer's purchased weight β , agent A 's purchased weight α , and the gain-from-trade Δ_a^t . In this setting, we have that

Theorem 4.1. Bounds on Buyer's Gain-from-Trade: *In the linear model setting, by knowing Δ_a^t and weights α, β , agent A can obtain bounds on the gain-from-trade of agent B given by:*

$$\left(\frac{1 - \sqrt{\Delta_a^t}(1 - \alpha)}{(1 - \beta) + \sqrt{\Delta_a^t}(1 - \alpha - \beta + 2\alpha\beta)} \right)^2 \leq \Delta_b^t \leq \left(\frac{1 + \sqrt{\Delta_a^t}(1 - \alpha)}{(1 - \beta) - \sqrt{\Delta_a^t}(1 - \alpha - \beta + 2\alpha\beta)} \right)^2.$$

Discussion. Theorem 4.1 states that by knowing information disclosed by the broker (including purchased weights α, β), the seller agent A can find the upper and lower bounds of the buyer's gain-from-trade, Δ_b^t . This information can then be used to estimate the value that the buyer places on the item. Using the Bayesian optimal-pricing mechanism, and taking into account the known probability distribution, the seller can estimate the price to determine their own virtual valuation.

Furthermore, the optimal scenario occurs when the seller agent values self-parameters exactly as the buyer does. In this case, based on Eq. (1), the broker will set the transfer payment as $\Delta P_{ab}^t = P_a^t - P_b^t = v_b(\hat{\theta}_a^t) - v_a(\hat{\theta}_a^t)$. Hence, the transfer payment is equal to the difference between the gain-from-trade of the two agents, where $\Delta P_{ab}^t = \Delta_a^t - \Delta_b^t$.

The proof for Theorem 4.1 is in the Appendix C. We summarize trading steps in a single round for this instantiation above (Algorithm 1), where two agents are training and trading parameters for the linear model setting. Here, agents A, B act as a seller and a buyer, respectively.

5 Convergence Analysis

Next, we study a theoretical analysis for the *effectiveness of buying parameters*. In particular, we are interested in understanding whether participating in the market leads to faster training convergence (in the worst-case scenario). We show this holds in a simplified setting where agent A always leads and never purchases parameters, while agent B always purchases from agent A . This asymmetry could be due to various reasons including a lack of training resources for agent B . Here we study a setting for general L -smooth functions, which is more practical, as the broker doesn't need to be knowledgeable about the true parameter θ^* . We assume the broker's loss is lower than agents' losses, in particular, $\hat{\mathcal{L}}_z(\theta) \leq \hat{\mathcal{L}}_a(\theta)$ and $\hat{\mathcal{L}}_z(\theta) \leq \hat{\mathcal{L}}_b(\theta), \forall \theta \in \mathbb{R}^d$. We take the gain-from-trade Δ_u^t by using the subtraction of empirical loss before and after a trade. We write it as

$$\Delta_u^t = \hat{\mathcal{L}}_z(\hat{\theta}_u^t) - \hat{\mathcal{L}}_z(\bar{\theta}_u^t), u \in \{a, b\}. \quad (4)$$

Theorem 5.1. *For all agents $u \in \{a, b, z\}$, let the loss function $\hat{\mathcal{L}}_u$ be L -smooth and let the samples on all agents be drawn from the same distribution \mathcal{D} . Let $\mathbf{E}_{\mathcal{D}}[\hat{\mathcal{L}}_u] = \mathcal{L}_u$, and $\Delta_b^t = \hat{\mathcal{L}}_z(\hat{\theta}_b^t) - \hat{\mathcal{L}}_z(\bar{\theta}_b^t)$. Let the algorithm run until round T with step size $\eta \in (0, \frac{1}{L})$, and let $\delta_b := \min_{t \in [T]} \mathbf{E}[\Delta_b^t]$ and $\bar{g}_b^2 := \min_{t \in [T]} \mathbf{E}[\|\nabla \hat{\mathcal{L}}_b(\theta_b^t)\|_2^2]$. Then we have the following results,*

- a) *(Always Trade) If $\Delta_b^t > 0, \forall t$, and agent B always buys (i.e. $I_b^t = 1, \forall t$). Then $T \geq \frac{2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*))}{\eta \epsilon^2 + 2\delta_b}$ implies $\bar{g}_b \leq \epsilon$.*
- b) *(Never Trade) If the agents never trade i.e. ($I_a^t = I_b^t = 0, \forall t$). Then $\bar{g}_b \leq \epsilon$ for $T \geq \frac{2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*))}{\eta \epsilon^2}$.*

Discussion. We show the convergence rate of always trade for agent B is $\mathcal{O}(1/(\epsilon^2 + \delta_b))$, while never trade is $\mathcal{O}(1/\epsilon^2)$. The difference between these two scenarios is due to δ_b — the minimal gain-from-trade over T runs. More gain-from-trade implies a smaller T (i.e. faster convergence) in the worst-case scenario. In addition, when δ_b is $\Omega(\epsilon)$, we can get much better convergence rate of $\mathcal{O}(1/\epsilon)$. Our results in this fundamental setting illustrate that participation in the market can lead to faster convergence when there exists gain-from-trade. The proof for Theorem 5.1 is in the Appendix D.1. In addition to this general setting, we provide convergence analysis for the linear model that we used as a case study in Sec. 4. This analysis also shows a better convergence rate when agents trade parameters. See Appendix D.2 for more details.

6 Experiments

We study the proposed framework empirically. Our goals are to validate (i) trading in the proposed framework results in improvements, (ii) these improvements persist even when trading subsets of parameters, (iii) these persist even when agents are trading different models for different tasks, (iv) trading and pricing are viable in competitive settings, and (v) understand the importance of key components of our framework, such as the need for alignment.

6.1 Collaborative Agents

We first conduct experiments in a collaborative setting, where there is no payment for buying parameters. Our goal is to validate the potential improvements from transactions independently of the pricing mechanism, which we explore in a later set of experiments.

6.1.1 Parameter Trading in Neural Networks

Setup. We use MNIST [33], CIFAR10 [34], and TinyImageNet [35] for training MLPs and ResNet20 [36]. Agents have imbalanced datasets where half of the classes contain only 10% of datapoints. Agents are limited to collecting a part of a dataset, making it difficult for them to achieve satisfactory performance without collaborating and trading parameters to leverage each other's strengths.

Models are trained from different random initializations and batch orders over 60 epochs. Agents trade entire parameter sets and join the market after five epochs. The broker discloses gain-from-trade to agents. Broker aligns parameters [16], then merge.

		Agent A	Agent B
		Testing Acc. (%)	Testing Acc. (%)
MNIST + MLP	out-of-market	68.50%	72.97%
	FedAvg	81.98%	81.98%
	w/o alignment	84.64%	84.64%
	w alignment	86.96%	86.96%
CIFAR10 + ResNet20	out-of-market	71.14%	70.56%
	FedAvg	70.35%	67.85%
	w/o alignment	78.31%	78.31%
	w alignment	79.90%	79.90%
TinyImageNet + ResNet20	out-of-market	21.67%	15.89%
	FedAvg	19.95%	19.33%
	w/o alignment	31.28%	31.30%
	w alignment	31.69%	31.82%

Table 1: Testing accuracies are reported for each combination of dataset and model. In *FedAvg*, there is no broker to assist agents in conducting transactions. The interpolated weight is determined solely based on the proportion of data assets. *w/o alignment* indicates that broker merges parameters via simple interpolation with the optimized purchased weight. In *w alignment*, the broker aligns parameters by applying [16] and then interpolates.

In addition to the approach that agents train models on their own, we include another baseline method, FedAvg [37], which assumes that there is no broker involved in a trade to help agents align parameters and optimize their purchased weights. In FedAvg, the interpolated weight is determined by the portion of data assets that an agent is endowed with, which is 0.5 in this setting.

Results. Table 1 shows the performance of two agents. We find that both agents are able to achieve improved performance by leveraging each other’s training expertise, compared to out-of-market agents. Specifically, training and trading ResNet20 with TinyImageNet resulted in improving accuracy by **+10.02%** and **+15.93%**, respectively. We measure two ways to merge parameters for buying: with and without model alignment. With model alignment, the broker is able to merge models more effectively, ultimately enhancing performance for both agents. In addition, compared to FedAvg method, results confirm the significance of having a trusted broker in parameter trading. Without an intermediary broker to facilitate the trade, the performance of purchased weights can be negatively impacted, as evidenced by the results of CIFAR10 + ResNet20 and TinyImageNet + ResNet20.

Finally, Figure 2 displays a comparison of testing loss for the MLP on MNIST. Our results demonstrate that trading parameters results in faster convergence compared to siloed training. Using alignment further helps, as expected.

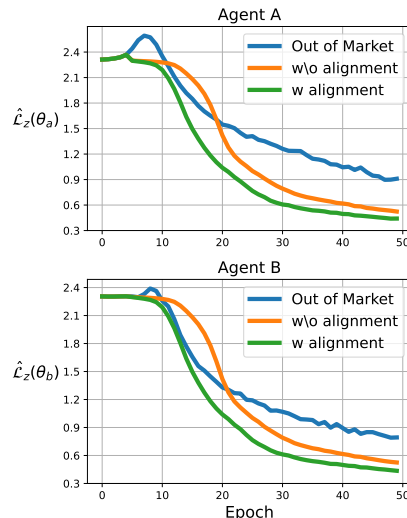


Figure 2: Testing loss converges the fastest by aligning and interpolating.

6.1.2 Parameter Subsets Trading in Neural Networks

Setup. Next, we explore the potential benefits of trading *subsets* of parameters. This scenario may take place when agents are interested in only certain components or are restricted by trading budgets. We use the same data endowment as in the preceding configuration. We train two 5-layer MLPs on MNIST and align parameters in each layer to trade.

Results. Table 2 displays the results by trading parameters from different layers. As expected, trading the entire model gives optimal performance (the last row). Trading subsets is helpful but suboptimal. We observe that purchasing parameters from shallow layers, close to the input layers,

	Agent A	Agent B
	Testing Acc. (%)	Testing Acc. (%)
out-of-market	71.29%	72.54%
layers {3, 4}	66.28%	71.98%
layers {2, 3, 4}	70.73%	73.36%
layers {0, 1, 2}	74.76%	74.16%
layers {0, 1}	78.86%	79.82%
layers {0, 1, 2, 3, 4}	86.96%	86.96%

Table 2: Testing accuracies for trading parameters from different layers in 5-layer MLPs on MNIST.




Figure 3: **Both agents earn improvements mutually through trading in the market.** We visualize trading results of two linear regression models with a synthesized dataset. Leftmost: trading logs over 100 runs. Second from left: The ratio of squared parameter estimation error between agent A and agent B. A red line represents tied performance. Above the red line, agent B leads, and vice versa. Second from right & rightmost: agent A’s and agent B’s learning curve compared to out-of-market agents. Market usage thus produces performance improvements.

offers more benefits than trading with deeper layers. This provides trading guidance for budget-limited agents.

6.1.3 The Effectiveness of Buying Parameters in Controlled Settings

Setup. We use a synthetic dataset to study trading in a fully-controlled environment. We use two linear regression models with dimension $d = 1000$. Agent A has $n_a = 500$ datapoints in dataset (X_a, Y_a) and B has $n_b = 800$ datapoints in dataset (X_b, Y_b) , but the latter’s labels are affected by zero-mean Gaussian noise ($\sigma^2 = 0.5$). We assume that the broker knows the true parameter θ^* and obtains (X_z, Y_z) and has access to $n_z = 10,000$ datapoints. Both agents start learning function f_a, f_b with the same initialized weight θ^0 . We compare the results over 100 runs with agents who obtain the same data endowment but do not participate in the market.

Results. The leftmost Figure 3 displays the trading log over 100 runs. Green dots and red dots show whether an agent purchases the other’s parameters in a specific run. Next, we show the ratio of squared parameter estimation error between agents. If the ratio is larger than 1 (red dashed line), agent B leads the market. We see that agent A leads the market in the first half of runs, making agent B continue buying parameters from agent A. At the end of a few runs, agent B turns to the lead.

The rightmost plots show convergence. If an agent is involved in the market and trades with the other, the convergence rate is faster compared to never trading. This study demonstrates agents’ behaviors in the trading runs and validates the effectiveness of buying parameters, leading to more efficient model training. We compute the empirical testing loss at the end. Compared to out-of-market agents, we find that agent A and agent B are able to improve testing loss by **42.84%** and **23.88%**, respectively.

6.1.4 Trading Parameters of Models with Different Purposes

Setup. Next, we validate whether trading makes sense even if agents are training *different models for different purposes*. This scenario is more realistic as it simulates situations where organizations in the market are working on different but potentially related tasks. We model this setting by sweeping the distance between the true parameters θ_a^*, θ_b^* of two linear regression models to observe how it impacts the benefits of trading.

Results. We record the benefits from trading when compared to an agent who does not participate but has the same data as agent




Figure 4: Trading with parameters that are from related tasks is possible.

A. We measure the relative improvement in empirical testing loss. The results are shown in Figure 4, which indicates that even though the two agents are not training on the same task, agent A is still able to benefit from trading. Note that the gain exists even when the tasks are quite different (i.e. large distance $\|\theta_a^* - \theta_b^*\|_2$) but is strongest as the tasks are most closely related.

6.2 Competitive Agents

Finally, we study agents engaging in a competitive scenario. In this case, the transactions require pricing. We validate the proposed pricing mechanism.

6.2.1 The Effectiveness of Bayesian Optimal Pricing Mechanism

Setup. We reuse the synthesized dataset from the collaborative experiment. We set the buyer’s valuation to gain-from-trade and estimate the seller’s virtual valuation by the lower bound that we can find.

Results. Market prices negotiated by the broker over 100 trading runs are displayed in Figure 5. Based on Eq. 1, the market price is determined by the average of the buyer’s valuation and the seller’s virtual valuation. To demonstrate market efficiency, we also study the scenario where the seller sets the valuation to be exactly the same as the buyer’s. As shown in Figure 3 (second from left: performance ratio), agent A is initially leading in the first half of the runs, resulting in a higher price for their parameters. However, at the end of a few runs, agent B takes the lead, causing their parameter price to increase while opponent A ’ goes down. It is important to note that the gap between the estimated price (the blue line) and the optimal price (the orange line) can be reduced with more information learned through the market, such as via historical transactions. Besides, the resulting negotiated price creates a discrepancy with the price in the optimal scenario where both parties report their valuations truthfully, *highlighting the significance of revealing accurate parameter values and justifying the need for incentives*. At last, this study illustrates the feasibility of using the Bayesian optimal pricing mechanism to assist seller agents in monetizing parameters with limited information to make a trade.




Figure 5: We visualize parameter market price negotiated by broker. We can see in the first half of runs, since agent A ’s performance dominates the market, making agent A ’s parameter more valuable compared to opponent agent B . Note that, if there is no trade, market price remains the same as historical price.

7 Conclusion

In this paper, we introduced a framework for parameter markets that can serve to reduce the heavy costs of large-scale model training. Borrowing from economic principles, we provided a set of mechanisms that enable the functioning of such markets. Theoretically, for simple settings, we analyzed market efficiency and proved that agents can gain from participating. We empirically validated the effectiveness of parameter markets under collaborative and competitive settings and demonstrated when participants in the market earn mutual benefits through market usage.

8 Acknowledgments

We would like to express our gratitude to the Wisconsin Alumni Research Foundation (WARF) for supporting this work. Additionally, we would like to thank all the reviewers for their valuable comments and constructive feedback on our work.

References

- [1] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [2] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020.
- [3] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- [4] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- [5] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [6] Or Sharir, Barak Peleg, and Yoav Shoham. The cost of training nlp models: A concise overview. *arXiv preprint arXiv:2004.08900*, 2020.
- [7] Lie He, An Bian, and Martin Jaggi. Cola: Decentralized linear learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- [8] Hanlin Tang, Xiangru Lian, Ming Yan, Ce Zhang, and Ji Liu. d^2 : Decentralized training over decentralized data. In *International Conference on Machine Learning*, pages 4848–4856. PMLR, 2018.
- [9] Qinyi Luo, Jinkun Lin, Youwei Zhuo, and Xuehai Qian. Hop: Heterogeneity-aware decentralized training. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 893–907, 2019.
- [10] Binhang Yuan, Yongjun He, Jared Quincy Davis, Tianyi Zhang, Tri Dao, Beidi Chen, Percy Liang, Christopher Re, and Ce Zhang. Decentralized training of foundation models in heterogeneous environments. *arXiv preprint arXiv:2206.01288*, 2022.
- [11] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018.
- [12] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pages 3259–3269. PMLR, 2020.
- [13] Norman Tatro, Pin-Yu Chen, Payel Das, Igor Melnyk, Prasanna Sattigeri, and Rongjie Lai. Optimizing mode connectivity via neuron alignment. *Advances in Neural Information Processing Systems*, 33:15300–15311, 2020.
- [14] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Dilan Gorur, Razvan Pascanu, and Hassan Ghasemzadeh. Linear mode connectivity in multitask and continual learning. *arXiv preprint arXiv:2010.04495*, 2020.
- [15] Rahim Entezari, Hanie Sedghi, Olga Saukh, and Behnam Neyshabur. The role of permutation invariance in linear mode connectivity of neural networks. *arXiv preprint arXiv:2110.06296*, 2021.
- [16] Samuel K Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. Git re-basin: Merging models modulo permutation symmetries. *arXiv preprint arXiv:2209.04836*, 2022.

- [17] Keller Jordan, Hanie Sedghi, Olga Saukh, Rahim Entezari, and Behnam Neyshabur. Repair: Renormalizing permuted activations for interpolation repair. *arXiv preprint arXiv:2211.08403*, 2022.
- [18] Min Du, Ruoxi Jia, and Dawn Song. Robust anomaly detection and backdoor attack detection via differential privacy. *arXiv preprint arXiv:1911.07116*, 2019.
- [19] Xudong Zhu and Hui Li. Privacy-preserving decentralized federated deep learning. In *ACM Turing Award Celebration Conference-China (ACM TURC 2021)*, pages 33–38, 2021.
- [20] Dario Pasquini, Mathilde Raynal, and Carmela Troncoso. On the privacy of decentralized machine learning. *arXiv preprint arXiv:2205.08443*, 2022.
- [21] Qiyang Song, Jiahao Cao, Kun Sun, Qi Li, and Ke Xu. Try before you buy: Privacy-preserving data evaluation on cloud-based machine learning data marketplace. In *Annual Computer Security Applications Conference*, pages 260–272, 2021.
- [22] Santiago Andrés Azcoitia and Nikolaos Laouraris. Try before you buy: A practical data purchasing algorithm for real-world data marketplaces. In *Proceedings of the 1st International Workshop on Data Economy*, pages 27–33, 2022.
- [23] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [24] John F Nash Jr. The bargaining problem. *Econometrica: Journal of the econometric society*, pages 155–162, 1950.
- [25] Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In *International Conference on Machine Learning*, pages 2242–2251. PMLR, 2019.
- [26] Anish Agarwal, Munther Dahleh, and Tuhin Sarkar. A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 701–726, 2019.
- [27] Raul Castro Fernandez, Pranav Subramaniam, and Michael J Franklin. Data market platforms: Trading data assets to solve data problems. *arXiv preprint arXiv:2002.01047*, 2020.
- [28] Junjie Chen, Minming Li, and Haifeng Xu. Selling data to a machine learner: Pricing via costly signaling. In *International Conference on Machine Learning*, pages 3336–3359. PMLR, 2022.
- [29] Meena Jagadeesan, Michael I Jordan, and Nika Haghtalab. Competition, alignment, and equilibria in digital marketplaces. *arXiv preprint arXiv:2208.14423*, 2022.
- [30] Lingjiao Chen, Paraschos Koutris, and Arun Kumar. Towards model-based pricing for machine learning in a data marketplace. In *Proceedings of the 2019 International Conference on Management of Data*, pages 1535–1552, 2019.
- [31] Jinfei Liu, Jian Lou, Junxu Liu, Li Xiong, Jian Pei, and Jimeng Sun. Dealer: an end-to-end model marketplace with differential privacy. 2021.
- [32] Charles W Cobb and Paul H Douglas. A theory of production. 1928.
- [33] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [34] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- [35] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- [36] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [37] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.

- [38] ABDULLAH YAVAŞ. Economics of brokerage: an overview. *Journal of Real Estate Literature*, 2(2):169–195, 1994.
- [39] Simran Khanuja, Melvin Johnson, and Partha Talukdar. Mergedistill: Merging pre-trained language models using distillation. *arXiv preprint arXiv:2106.02834*, 2021.
- [40] Sebastian Lawrenz, Priyanka Sharma, and Andreas Rausch. Blockchain technology as an approach for data marketplaces. In *Proceedings of the 2019 international conference on blockchain technology*, pages 55–59, 2019.
- [41] Franziska Boenisch. A systematic review on model watermarking for neural networks. *Frontiers in big Data*, 4:729663, 2021.
- [42] Reza Yazdani Aminabadi, Samyam Rajbhandari, Ammar Ahmad Awan, Cheng Li, Du Li, Elton Zheng, Olatunji Ruwase, Shaden Smith, Minjia Zhang, Jeff Rasley, et al. Deepspeed-inference: enabling efficient inference of transformer models at unprecedented scale. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15. IEEE, 2022.
- [43] Jannik Kossen, Sebastian Farquhar, Yarin Gal, and Tom Rainforth. Active testing: Sample-efficient model evaluation. In *International Conference on Machine Learning*, pages 5753–5763. PMLR, 2021.
- [44] Ante Ćustić, Vladyslav Sokol, Abraham P Punnen, and Binay Bhattacharya. The bilinear assignment problem: complexity and polynomially solvable special cases. *Mathematical programming*, 166:185–205, 2017.

The appendix is organized as follows. In Appendix A, we show steps for broker to solve revenue maximization problem in a trade (Sec. 3.4). Next, we provide details for the instantiation in Sec. 4. First, in Appendix B, we offer proof for agents to bound performance ratio under different trading decisions. Then, in Appendix C, we provide proof of Theorem 4.1 for agents to bound the other agent's gain-from-trade. We generalize instantiation to a broader setting for L-smooth function and study its convergence analysis in Appendix D.1. In addition to L-smooth function, we use linear models as a case study and provide its convergence analysis in Appendix D.2. In Appendix E, we place experimental details regarding data endowments and training settings. Then, in Appendix F, we offer more findings as trading guidance investigated by different trading scenarios. At last, we discuss limitations, potential concerns, computational needs for broker, and agents' common priors in Appendix G.

A Pricing Mechanism

In this section, we offer steps to maximize revenue of both agents. Recall that we define agent u 's valuation as $v_u(\theta)$ and the market price of agent u 's parameters at time t as P_u^t . We set agents' revenue as follows,

$$U_a(P_a^t, P_b^t) := (P_a^t - v_a(\hat{\theta}_a^t)) + (v_a(\hat{\theta}_b^t) - P_b^t) \quad U_b(P_b^t, P_a^t) := (P_b^t - v_b(\hat{\theta}_b^t)) + (v_b(\hat{\theta}_a^t) - P_a^t)$$

We formulate the revenue maximization problem accordingly,

$$\begin{aligned} & \operatorname{argmax}_{P_a^t, P_b^t} U_a(P_a^t, P_b^t) \times U_b(P_b^t, P_a^t) \\ \text{s.t.} \quad & P_a^t \in [v_a(\hat{\theta}_a^t), v_b(\hat{\theta}_a^t)], \quad P_b^t \in [v_b(\hat{\theta}_b^t), v_b(\hat{\theta}_a^t)] \end{aligned}$$

Let the difference in price that broker finds as $\Delta P_{ab}^t := P_a^t - P_b^t$. Then we can have,

$$\begin{aligned} & \operatorname{argmax}_{P_a^t, P_b^t} \left[(P_a^t - v_a(\hat{\theta}_a^t)) + (v_a(\hat{\theta}_b^t) - P_b^t) \right] \times \left[(P_b^t - v_b(\hat{\theta}_b^t)) + (v_b(\hat{\theta}_a^t) - P_a^t) \right] \\ & \Rightarrow \operatorname{argmax}_{P_a^t, P_b^t} (\Delta P_{ab}^t - v_a(\hat{\theta}_a^t) + v_a(\hat{\theta}_b^t)) \times (-\Delta P_{ab}^t - v_b(\hat{\theta}_b^t) + v_b(\hat{\theta}_a^t)) \\ & = \operatorname{argmax}_{P_a^t, P_b^t} \left(-(\Delta P_{ab}^t)^2 - v_b(\hat{\theta}_b^t) \cdot \Delta P_{ab}^t + v_b(\hat{\theta}_a^t) \cdot \Delta P_{ab}^t + v_a(\hat{\theta}_a^t) \cdot \Delta P_{ab}^t + v_a(\hat{\theta}_a^t) \cdot v_b(\hat{\theta}_b^t) \right. \\ & \quad \left. - v_a(\hat{\theta}_a^t) \cdot v_b(\hat{\theta}_a^t) - v_a(\hat{\theta}_b^t) \cdot \Delta P_{ab}^t - v_a(\hat{\theta}_b^t) \cdot v_b(\hat{\theta}_b^t) + v_a(\hat{\theta}_b^t) \cdot v_b(\hat{\theta}_a^t) \right) \end{aligned}$$

Taking the derivative and set equation to zero,

$$-2\Delta P_{ab}^t - v_b(\hat{\theta}_b^t) + v_b(\hat{\theta}_a^t) + v_a(\hat{\theta}_a^t) - v_a(\hat{\theta}_b^t) = 0$$

Then we can find that the ΔP_{ab}^t by

$$\Delta P_{ab}^t = P_a^t - P_b^t = \frac{1}{2} \left(v_b(\hat{\theta}_a^t) + v_a(\hat{\theta}_a^t) - v_a(\hat{\theta}_b^t) - v_b(\hat{\theta}_b^t) \right)$$

The resulting price difference represents the transferred payment between agents.

B Performance Ratio

We provide a concrete example of how agents value parameters in the linear model in Sec. 4. To approach Theorem 4.1, we start by finding bounds on the performance ratio between agents under different trading decisions. Here, the performance ratio in the market at time t is determined by the following, and we use agent A as an example to bound the ratio between agent A and his opponent.

$$\frac{\|\theta_b^t - \theta^*\|_2^2}{\|\theta_a^t - \theta^*\|_2^2}$$

Once the ratio is greater than 1, agent A is the lead in the market. Note that θ_a^t and θ_b^t are the final parameters based on agents' decisions, and θ^* is the true parameter known by broker. In a two-agent market, There are four different scenarios to analyze.

$$\{\theta_a^t, \theta_b^t\} = \begin{cases} \hat{\theta}_a^t, \hat{\theta}_b^t & \text{if agent } A \text{ doesn't buy and doesn't sell parameters,} \\ \bar{\theta}_a^t, \hat{\theta}_b^t & \text{if agent } A \text{ buys but doesn't sell parameters,} \\ \hat{\theta}_a^t, \bar{\theta}_b^t & \text{if agent } A \text{ doesn't buy but sell parameters,} \\ \bar{\theta}_a^t, \bar{\theta}_b^t & \text{if agent } A \text{ buys and sells parameters.} \end{cases}$$

Recall that the purchased parameters are defined as

$$\begin{aligned} \bar{\theta}_a^t &:= (1 - \alpha)\hat{\theta}_a^t + \alpha\hat{\theta}_b^t, & \alpha \in (0, 1] \\ \bar{\theta}_b^t &:= (1 - \beta)\hat{\theta}_b^t + \beta\hat{\theta}_a^t, & \beta \in (0, 1] \end{aligned}$$

In the instantiation, recall that we take the gain-from-trade revealed by broker for agent A to be

$$\Delta_a^t := \frac{\|\hat{\theta}_a^t - \theta^*\|_2^2}{\|\bar{\theta}_a^t - \theta^*\|_2^2}$$

Lemma B.1. Agent A: (No-Buy, No-Sell) *If agent A doesn't buy parameters from agent B and doesn't sell parameters to agent B, then knowing the purchased weight α and gain-from-trade Δ_a^t , the ratio of parameter estimation error between agent A and agent B is bounded by:*

$$\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right)^2 \leq \frac{\|\hat{\theta}_b^t - \theta^*\|_2^2}{\|\hat{\theta}_a^t - \theta^*\|_2^2} \leq \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1 - \alpha}{\alpha}\right)^2$$

Lemma B.2. Agent A: (Buy, No-Sell) *If agent A buys parameters from agent B and doesn't sell parameters to agent B, then knowing the purchased weight α and gain-from-trade Δ_a^t , the ratio of parameter estimation error between agent A and agent B is bounded by:*

$$\Delta_a^t \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right)^2 \leq \frac{\|\hat{\theta}_b^t - \theta^*\|_2^2}{\|\hat{\theta}_a^t - \theta^*\|_2^2} \leq \Delta_a^t \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1 - \alpha}{\alpha}\right)^2$$

Proof.

$$\begin{aligned} \|\hat{\theta}_b^t - \theta^*\|_2^2 &= \left\| \frac{\bar{\theta}_a^t - (1 - \alpha)\hat{\theta}_a^t}{\alpha} - \theta^* \right\|_2^2 \\ &= \left\| \frac{1}{\alpha}(\bar{\theta}_a^t - \theta^*) - \frac{1 - \alpha}{\alpha}(\hat{\theta}_a^t - \theta^*) \right\|_2^2 \\ &= \frac{1}{\alpha^2} \|\bar{\theta}_a^t - \theta^*\|_2^2 + \left(\frac{1 - \alpha}{\alpha}\right)^2 \|\hat{\theta}_a^t - \theta^*\|_2^2 - \frac{2(1 - \alpha)}{\alpha^2} \langle \bar{\theta}_a^t - \theta^*, \hat{\theta}_a^t - \theta^* \rangle \\ &\leq \frac{1}{\alpha^2} \|\bar{\theta}_a^t - \theta^*\|_2^2 + \left(\frac{1 - \alpha}{\alpha}\right)^2 \|\hat{\theta}_a^t - \theta^*\|_2^2 + \frac{2(1 - \alpha)}{\alpha^2} \|\bar{\theta}_a^t - \theta^*\|_2 \|\hat{\theta}_a^t - \theta^*\|_2 \\ &= \frac{1}{\Delta_a^t \alpha^2} \|\hat{\theta}_a^t - \theta^*\|_2^2 + \left(\frac{1 - \alpha}{\alpha}\right)^2 \|\hat{\theta}_a^t - \theta^*\|_2^2 + \frac{2(1 - \alpha)}{\sqrt{\Delta_a^t} \alpha^2} \|\hat{\theta}_a^t - \theta^*\|_2^2 \\ &= \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1 - \alpha}{\alpha}\right)^2 \|\hat{\theta}_a^t - \theta^*\|_2^2 \end{aligned}$$

Hence, we can have an upper bound

$$\frac{\|\hat{\theta}_b^t - \theta^*\|_2^2}{\|\hat{\theta}_a^t - \theta^*\|_2^2} \leq \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1 - \alpha}{\alpha}\right)^2$$

For the lower bound,

$$\begin{aligned}
\|\dot{\theta}_b^t - \theta^*\|_2^2 &= \left\| \frac{\bar{\theta}_a^t - (1-\alpha)\dot{\theta}_a^t}{\alpha} - \theta^* \right\|_2^2 \\
&= \left\| \frac{1}{\alpha}(\bar{\theta}_a^t - \theta^*) - \frac{1-\alpha}{\alpha}(\dot{\theta}_a^t - \theta^*) \right\|_2^2 \\
&= \frac{1}{\alpha^2} \|\bar{\theta}_a^t - \theta^*\|_2^2 + \left(\frac{1-\alpha}{\alpha}\right)^2 \|\dot{\theta}_a^t - \theta^*\|_2^2 - \frac{2(1-\alpha)}{\alpha^2} \langle \bar{\theta}_a^t - \theta^*, \dot{\theta}_a^t - \theta^* \rangle \\
&\geq \frac{1}{\alpha^2} \|\bar{\theta}_a^t - \theta^*\|_2^2 + \left(\frac{1-\alpha}{\alpha}\right)^2 \|\dot{\theta}_a^t - \theta^*\|_2^2 - \frac{2(1-\alpha)}{\alpha^2} \|\bar{\theta}_a^t - \theta^*\|_2 \|\dot{\theta}_a^t - \theta^*\|_2 \\
&= \frac{1}{\Delta_a^t \alpha^2} \|\dot{\theta}_a^t - \theta^*\|_2^2 + \left(\frac{1-\alpha}{\alpha}\right)^2 \|\dot{\theta}_a^t - \theta^*\|_2^2 - \frac{2(1-\alpha)}{\sqrt{\Delta_a^t} \alpha^2} \|\dot{\theta}_a^t - \theta^*\|_2^2 \\
&= \left(\frac{1}{\alpha \sqrt{\Delta_a^t}} - \frac{1-\alpha}{\alpha}\right)^2 \|\dot{\theta}_a^t - \theta^*\|_2^2
\end{aligned}$$

Hence, we can have a lower bound satisfied by

$$\frac{\|\dot{\theta}_b^t - \theta^*\|_2^2}{\|\dot{\theta}_a^t - \theta^*\|_2^2} \geq \left(\frac{1}{\alpha \sqrt{\Delta_a^t}} - \frac{1-\alpha}{\alpha}\right)^2$$

Therefore, by knowing gain-from-trade Δ_a^t and purchased weight α from the broker, in {No-Buy, No-Sell} scenario, agent A can bound the performance ratio with the final parameter $\dot{\theta}_a^t$ and $\dot{\theta}_b^t$. We can also have bounds for another scenario {Buy, No-Sell} by multiplying Δ_a^t to the inequality. \square

Another case in the market is agent B wishes to buy parameters from agent A . In this case, agent A will receive a quotation request from agent B with his purchased weight β . Then, agent A is able to bound the performance ratio.

Lemma B.3. Agent A: (No-Buy, Sell) *If agent A doesn't buy parameters from agent B but sells parameters to agent B , then knowing purchased weights α, β and gain-from-trade Δ_a^t , the ratio of parameter estimation error between agent A and agent B is bounded by:*

$$\left[(1-\beta)\left(\frac{1}{\alpha \sqrt{\Delta_a^t}} - \frac{1-\alpha}{\alpha}\right) - \beta\right]^2 \leq \frac{\|\bar{\theta}_b^t - \theta^*\|_2^2}{\|\dot{\theta}_a^t - \theta^*\|_2^2} \leq \left[(1-\beta)\left(\frac{1}{\alpha \sqrt{\Delta_a^t}} + \frac{1-\alpha}{\alpha}\right) + \beta\right]^2$$

Lemma B.4. Agent A: (Buy, Sell) *If agent A buys parameters from agent B and sells parameters to agent B , then knowing purchased weights α, β and gain-from-trade Δ_a^t , the ratio of parameter estimation error between agent A and agent B is bounded by:*

$$\Delta_a^t \left[(1-\beta)\left(\frac{1}{\alpha \sqrt{\Delta_a^t}} - \frac{1-\alpha}{\alpha}\right) - \beta\right]^2 \leq \frac{\|\bar{\theta}_b^t - \theta^*\|_2^2}{\|\dot{\theta}_a^t - \theta^*\|_2^2} \leq \Delta_a^t \left[(1-\beta)\left(\frac{1}{\alpha \sqrt{\Delta_a^t}} + \frac{1-\alpha}{\alpha}\right) + \beta\right]^2$$

Proof.

$$\begin{aligned}
\|\bar{\theta}_b^t - \theta^*\|_2^2 &= \|(1-\beta)\dot{\theta}_b^t + \beta\dot{\theta}_a^t - \theta^*\|_2^2 \\
&= \|(1-\beta)(\dot{\theta}_b^t - \theta^*) + \beta(\dot{\theta}_a^t - \theta^*)\|_2^2 \\
&= (1-\beta)^2 \|\dot{\theta}_b^t - \theta^*\|_2^2 + \beta^2 \|\dot{\theta}_a^t - \theta^*\|_2^2 + 2(1-\beta)\beta \langle \dot{\theta}_b^t - \theta^*, \dot{\theta}_a^t - \theta^* \rangle \\
&\leq (1-\beta)^2 \|\dot{\theta}_b^t - \theta^*\|_2^2 + \beta^2 \|\dot{\theta}_a^t - \theta^*\|_2^2 + 2(1-\beta)\beta \|\dot{\theta}_b^t - \theta^*\|_2 \|\dot{\theta}_a^t - \theta^*\|_2 \\
&\leq (1-\beta)^2 \left(\frac{1}{\alpha \sqrt{\Delta_a^t}} + \frac{1-\alpha}{\alpha}\right)^2 \|\dot{\theta}_a^t - \theta^*\|_2^2 + \beta^2 \|\dot{\theta}_a^t - \theta^*\|_2^2 \\
&\quad + 2(1-\beta)\beta \left(\frac{1}{\alpha \sqrt{\Delta_a^t}} + \frac{1-\alpha}{\alpha}\right) \|\dot{\theta}_a^t - \theta^*\|_2^2 \tag{Lemma B.1} \\
&= \left[(1-\beta)\left(\frac{1}{\alpha \sqrt{\Delta_a^t}} + \frac{1-\alpha}{\alpha}\right) + \beta\right]^2 \|\dot{\theta}_a^t - \theta^*\|_2^2
\end{aligned}$$

Hence, we can have an upper bound

$$\frac{\|\bar{\theta}_b^t - \theta^*\|_2^2}{\|\dot{\theta}_a^t - \theta^*\|_2^2} \leq [(1 - \beta)\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1 - \alpha}{\alpha}\right) + \beta]^2$$

For the lower bound,

$$\begin{aligned} \|\bar{\theta}_b^t - \theta^*\|_2^2 &= \|(1 - \beta)\dot{\theta}_b^t + \beta\dot{\theta}_a^t - \theta^*\|_2^2 \\ &= \|(1 - \beta)(\dot{\theta}_b^t - \theta^*) + \beta(\dot{\theta}_a^t - \theta^*)\|_2^2 \\ &= (1 - \beta)^2\|\dot{\theta}_b^t - \theta^*\|_2^2 + \beta^2\|\dot{\theta}_a^t - \theta^*\|_2^2 + 2(1 - \beta)\beta\langle \dot{\theta}_b^t - \theta^*, \dot{\theta}_a^t - \theta^* \rangle \\ &\geq (1 - \beta)^2\|\dot{\theta}_b^t - \theta^*\|_2^2 + \beta^2\|\dot{\theta}_a^t - \theta^*\|_2^2 - 2(1 - \beta)\beta\|\dot{\theta}_b^t - \theta^*\|_2\|\dot{\theta}_a^t - \theta^*\|_2 \\ &\geq (1 - \beta)^2\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right)^2\|\dot{\theta}_a^t - \theta^*\|_2^2 + \beta^2\|\dot{\theta}_a^t - \theta^*\|_2^2 \\ &\quad - 2(1 - \beta)\beta\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right)\|\dot{\theta}_a^t - \theta^*\|_2^2 \tag{Lemma B.1} \\ &= [(1 - \beta)\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right) - \beta]^2\|\dot{\theta}_a^t - \theta^*\|_2^2 \end{aligned}$$

Hence, we can have a lower bound

$$\frac{\|\bar{\theta}_b^t - \theta^*\|_2^2}{\|\dot{\theta}_a^t - \theta^*\|_2^2} \geq [(1 - \beta)\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right) - \beta]^2$$

Therefore, by knowing gain-from-trade Δ_a^t and purchased weight α, β , in {No-Buy, Sell} scenario, agent A can bound the performance ratio with the final parameter $\dot{\theta}_a^t$ and $\bar{\theta}_b^t$. We can also have bounds for another scenario {Buy, Sell} by multiplying Δ_a^t to the inequality. \square

C Buyer's Gain-from-Trade

Next, we use Lemma B.1 and Lemma B.3 to approach Theorem 4.1. Recall that agent B 's gain-from-trade Δ_b^t at the time t is defined as,

$$\Delta_b^t := \frac{\|\dot{\theta}_b^t - \theta^*\|_2^2}{\|\dot{\theta}_b^t - \theta^*\|_2^2}$$

Theorem C.1. Bounds on Buyer's Gain-from-Trade: By knowing purchased weights α, β and Δ_a^t , agent A can obtain bounds on the gain-from-trade of agent B given by:

$$\left(\frac{1 - \sqrt{\Delta_a^t}(1 - \alpha)}{(1 - \beta) + \sqrt{\Delta_a^t}(1 - \alpha - \beta + 2\alpha\beta)}\right)^2 \leq \Delta_b^t \leq \left(\frac{1 + \sqrt{\Delta_a^t}(1 - \alpha)}{(1 - \beta) - \sqrt{\Delta_a^t}(1 - \alpha - \beta + 2\alpha\beta)}\right)^2.$$

Proof. Based on Lemma B.1, we have

$$\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right)^2\|\dot{\theta}_a^t - \theta^*\|_2^2 \leq \|\dot{\theta}_b^t - \theta^*\|_2^2 \leq \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1 - \alpha}{\alpha}\right)^2\|\dot{\theta}_a^t - \theta^*\|_2^2$$

Dividing by $\|\bar{\theta}_b^t - \theta^*\|_2^2$, we can have

$$\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right)^2 \frac{\|\dot{\theta}_a^t - \theta^*\|_2^2}{\|\dot{\theta}_b^t - \theta^*\|_2^2} \leq \Delta_b^t \leq \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1 - \alpha}{\alpha}\right)^2 \frac{\|\dot{\theta}_a^t - \theta^*\|_2^2}{\|\dot{\theta}_b^t - \theta^*\|_2^2}$$

Since we know Lemma B.3 which gives us

$$[(1 - \beta)\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1 - \alpha}{\alpha}\right) - \beta]^2 \leq \frac{\|\bar{\theta}_b^t - \theta^*\|_2^2}{\|\dot{\theta}_a^t - \theta^*\|_2^2} \leq [(1 - \beta)\left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1 - \alpha}{\alpha}\right) + \beta]^2$$

Inverting the fraction, we can have

$$\left(\frac{\alpha\sqrt{\Delta_a^t}}{(1-\beta) + \sqrt{\Delta_a^t}(1-\alpha-\beta+2\alpha\beta)} \right)^2 \leq \frac{\|\dot{\theta}_a^t - \theta^*\|_2^2}{\|\dot{\theta}_b^t - \theta^*\|_2^2} \leq \left(\frac{\alpha\sqrt{\Delta_a^t}}{(1-\beta) - \sqrt{\Delta_a^t}(1-\alpha-\beta+2\alpha\beta)} \right)^2$$

Therefore, using Lemma B.1, the upper bound of Δ_b^t can be found by

$$\begin{aligned} \Delta_b^t &\leq \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} + \frac{1-\alpha}{\alpha} \right)^2 \left(\frac{\alpha\sqrt{\Delta_a^t}}{(1-\beta) - \sqrt{\Delta_a^t}(1-\alpha-\beta+2\alpha\beta)} \right)^2 \\ &\leq \left(\frac{1 + \sqrt{\Delta_a^t}(1-\alpha)}{(1-\beta) - \sqrt{\Delta_a^t}(1-\alpha-\beta+2\alpha\beta)} \right)^2 \end{aligned}$$

On the other side, we can have the lower bound as follows

$$\begin{aligned} \Delta_b^t &\geq \left(\frac{1}{\alpha\sqrt{\Delta_a^t}} - \frac{1-\alpha}{\alpha} \right)^2 \left(\frac{\alpha\sqrt{\Delta_a^t}}{(1-\beta) + \sqrt{\Delta_a^t}(1-\alpha-\beta+2\alpha\beta)} \right)^2 \\ &\geq \left(\frac{1 - \sqrt{\Delta_a^t}(1-\alpha)}{(1-\beta) + \sqrt{\Delta_a^t}(1-\alpha-\beta+2\alpha\beta)} \right)^2 \end{aligned}$$

Theorem 4.1 asserts that if the broker discloses certain information; purchased weights α, β and gain-from-trade Δ_a^t , agent A as a seller can determine the maximum and minimum values of buyer's gain-from-trade, Δ_b^t . This knowledge can then be utilized to approximate buyer's valuation so that seller's virtual valuation can be set. \square

D Convergence Analysis with Trading

D.1 (General) L-smooth Functions

In Sec. 5, we study a broader setting for general L-smooth functions and offer its convergence analysis to validate the *effectiveness of buying parameters*. In the general case, the broker informs the gain-from-trade by using the subtraction of empirical loss before and after a trade. This can be more practical, as the broker doesn't need to be knowledgeable about the true parameter θ^* . We write the gain-from-trade as follows,

$$\Delta_u^t := \hat{\mathcal{L}}_z(\dot{\theta}_u^t) - \hat{\mathcal{L}}_z(\bar{\theta}_u^t)$$

Theorem D.1. *For all agents $u \in \{a, b, z\}$, let the loss function $\hat{\mathcal{L}}_u$ be L -smooth, and let the samples on all agents be drawn from the same distribution \mathcal{D} . Let $\mathbf{E}_{\mathcal{D}}[\hat{\mathcal{L}}_u] = \mathcal{L}_u$, and $\Delta_b^t = \hat{\mathcal{L}}_z(\dot{\theta}_b^t) - \hat{\mathcal{L}}_z(\bar{\theta}_b^t)$. Let the algorithm run until round T with step size $\eta \in (0, \frac{1}{L})$, and let $\delta_b := \min_{t \in [T]} \mathbf{E}[\Delta_b^t]$ and $\bar{g}_b^2 := \min_{t \in [T]} \mathbf{E}[\|\nabla \hat{\mathcal{L}}_b(\theta_b^t)\|_2^2]$. Then we have the following results,*

- a) *(Always Trade) If $\Delta_b^t > 0, \forall t$, and agent B always buys (i.e. $I_b^t = 1, \forall t$). Then $T \geq \frac{2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*))}{\eta\epsilon^2 + 2\delta_b}$ implies $\bar{g}_b \leq \epsilon$.*
- b) *(Never Trade) If the agents never trade i.e. ($I_a^t = I_b^t = 0, \forall t$). Then $\bar{g}_b \leq \epsilon$ for $T \geq \frac{2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*))}{\eta\epsilon^2}$.*

Proof. The proof for the never trade case follows from standard analysis of gradient descent for L-smooth functions (see the proof of Proposition D.2). Here we provide proof for the always trade case. In this case we have $\Delta_b^t > 0$ for all rounds, causing agent B always buy parameters. The population level loss function is defined as \mathcal{L} , it is the expectation of $\hat{\mathcal{L}}_u, u \in \{a, b, z\}$. Hence, the expectation of the gain-from-trade for agent B can be written as

$$\mathbf{E}[\Delta_b^t] = \mathbf{E}[\hat{\mathcal{L}}_b(\dot{\theta}_b^t)] - \mathbf{E}[\hat{\mathcal{L}}_b(\bar{\theta}_b^t)] = \mathcal{L}(\dot{\theta}_b^t) - \mathcal{L}(\bar{\theta}_b^t) \quad (5)$$

Since the loss function, $\hat{\mathcal{L}}$ is L-smooth, thus we have the following descent lemma (see Eq. Descent Lemma in the proof of Proposition D.2),

$$\mathcal{L}(\hat{\theta}_b^t) \leq \mathcal{L}(\theta_b^{t-1}) - \frac{\eta}{2} \mathbf{E}[\|\nabla \hat{\mathcal{L}}_b(\theta_b^{t-1})\|_2^2], \quad \forall t$$

Using Eq. 5, we can have $\mathcal{L}(\hat{\theta}_b^t) = \mathcal{L}(\bar{\theta}_b^t) + \mathbf{E}[\Delta_b^t]$, substituting it in the above descent lemma we get,

$$\mathcal{L}(\bar{\theta}_b^t) \leq \mathcal{L}(\theta_b^{t-1}) - \frac{\eta}{2} \mathbf{E}[\|\nabla \hat{\mathcal{L}}_b(\theta_b^{t-1})\|_2^2] - \mathbf{E}[\Delta_b^t], \quad \forall t \quad (6)$$

Note that, since agent B always purchases parameters, the final parameter θ_b^t is exactly as same as $\bar{\theta}_b^t$. Then we sum Eq. 6 over T rounds, we obtain

$$\mathcal{L}(\theta_b^T) = \mathcal{L}(\bar{\theta}_b^T) \leq \mathcal{L}(\theta_b^0) - \frac{\eta}{2} \sum_{i=1}^T \mathbf{E}[\|\nabla \hat{\mathcal{L}}_b(\theta_b^{i-1})\|_2^2] - \sum_{i=1}^T \mathbf{E}[\Delta_b^i]$$

Let $\delta_b := \min_t \mathbf{E}[\Delta_b^t], \forall t$ and $\bar{g}_b^2 := \min_t \mathbf{E}[\|\nabla \hat{\mathcal{L}}_b(\theta_b^t)\|_2^2]$. Then,

$$\begin{aligned} \mathcal{L}(\theta_b^*) &\leq \mathcal{L}(\theta_b^T) \leq \mathcal{L}(\theta_b^0) - \frac{\eta T}{2} \bar{g}_b^2 - T \delta_b \\ \bar{g}_b^2 &\leq \frac{2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*))}{\eta T} - \frac{2\delta_b}{\eta} \end{aligned}$$

Want the R.H.S. to be at most ϵ^2 ,

$$\begin{aligned} \frac{2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*))}{\eta T} - \frac{2\delta_b}{\eta} &\leq \epsilon^2 \\ 2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*)) - 2\delta_b T &\leq \eta \epsilon^2 T \\ T &\geq \frac{2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*))}{\eta \epsilon^2 + 2\delta_b} \end{aligned}$$

Leading to convergence rate of $\mathcal{O}(1/(\epsilon^2 + \delta_b))$. Additionally, when δ_b is $\Omega(\epsilon)$, then we get a much better convergence rate of $\mathcal{O}(1/\epsilon)$. \square

Proposition D.2. (Never trade, L-smooth Loss) Let the loss functions $\hat{\mathcal{L}}_u$ be L-smooth. Let the samples on all agents be drawn from the same distribution \mathcal{D} , and $\mathbf{E}_{\mathcal{D}}[\hat{\mathcal{L}}_u] = \mathcal{L}_u$. Let agent B never trade (i.e. $I_b^t = 0, \forall t$). Let the algorithm run until round T with step size $\eta \in (0, \frac{1}{L})$, and let $\bar{g}_b^2 := \min_{t \in [T]} \mathbf{E}[\|\nabla \hat{\mathcal{L}}_b(\theta_b^t)\|_2^2]$. Then $\bar{g}_b \leq \epsilon$, for

$$T \geq \frac{2(\mathcal{L}(\theta_b^0) - \mathcal{L}(\theta_b^*))}{\eta \epsilon^2}$$

Proof. The proof is a standard analysis of gradient descent for L-smooth functions. We provide the proof here for reference. Due to the L-smoothness of the loss functions we have,

$$\hat{\mathcal{L}}_z(\hat{\theta}_b^t) \leq \hat{\mathcal{L}}_z(\theta_b^{t-1}) + \langle \nabla \hat{\mathcal{L}}_z(\theta_b^{t-1}), \hat{\theta}_b^t - \theta_b^{t-1} \rangle + \frac{L}{2} \|\hat{\theta}_b^t - \theta_b^{t-1}\|_2^2$$

Since, $\hat{\theta}_b^t = \theta_b^{t-1} - \eta \nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})$,

$$\begin{aligned} \hat{\mathcal{L}}_z(\hat{\theta}_b^t) &\leq \hat{\mathcal{L}}_z(\theta_b^{t-1}) - \eta \|\nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})\|_2^2 + \frac{\eta^2 L}{2} \|\nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})\|_2^2 \\ &= \hat{\mathcal{L}}_z(\theta_b^{t-1}) - \left(\eta - \frac{\eta^2 L}{2} \right) \|\nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})\|_2^2 \end{aligned}$$

The constant $\eta - \frac{\eta^2 L}{2}$ is lower bounded by $\eta/2$ for $\eta \in (0, 1/L)$. Leading to the following descent lemma,

$$\hat{\mathcal{L}}_z(\hat{\theta}_b^t) \leq \hat{\mathcal{L}}_z(\theta_b^{t-1}) - \frac{\eta}{2} \|\nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})\|_2^2$$

Taking expectation over \mathcal{D} ,

$$\mathcal{L}(\hat{\theta}_b^t) \leq \mathcal{L}(\theta_b^{t-1}) - \frac{\eta}{2} \mathbf{E}[\|\nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})\|_2^2] \quad (\text{Descent Lemma})$$

Since agent B never trades, making $\theta_b^t = \hat{\theta}_b^t$,

$$\mathcal{L}(\theta_b^t) \leq \mathcal{L}(\theta_b^{t-1}) - \frac{\eta}{2} \mathbf{E}[\|\nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})\|_2^2]$$

Now, summing the above equation of T rounds, we can have

$$\begin{aligned} \mathcal{L}(\theta^*) &\leq \mathcal{L}(\theta_b^T) \leq \mathcal{L}(\theta_b^0) - \frac{\eta}{2} \sum_{t=1}^T \mathbf{E}[\|\nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})\|_2^2] \\ \sum_{t=1}^T \mathbf{E}[\|\nabla \hat{\mathcal{L}}_z(\theta_b^{t-1})\|_2^2] &\leq \frac{2(\mathcal{L}(\theta^0) - \mathcal{L}(\theta^*))}{\eta} \end{aligned}$$

Since $\bar{g}_b^2 := \min_{t \in [T]} \mathbf{E}[\|\nabla \hat{\mathcal{L}}_b(\theta_b^t)\|_2^2]$,

$$\bar{g}_b^2 \leq \frac{2(\mathcal{L}(\theta^0) - \mathcal{L}(\theta^*))}{\eta T}$$

Solving for $\bar{g}_b \leq \epsilon$ gives us,

$$T \geq \frac{2(\mathcal{L}(\theta^0) - \mathcal{L}(\theta^*))}{\eta \epsilon^2}$$

The convergence rate of *never trade* is $\mathcal{O}(1/\epsilon^2)$, which can be slower than *always trade* $\mathcal{O}(1/(\epsilon^2 + \delta_b))$ by the constant factor, gain-from-trade. \square

D.2 Linear Model Case Study

In addition to general case, we study the convergence for the instantiation in Sec. 4, where agents are trading parameters in linear models, and the gain-from-trade computed by the broker is defined as

$$\Delta_u^t := \frac{\|\hat{\theta}_u^t - \theta^*\|_2^2}{\|\hat{\theta}_u^t - \theta^*\|_2^2}, \quad u \in \{a, b\}$$

In this setting, we write the empirical loss function as

$$\hat{\mathcal{L}}_u(\theta) = \|X_u \theta - Y_u\|_2^2$$

Recall that the updated rules are

$$\begin{aligned} \hat{\theta}_a^t &= \theta_a^{t-1} - \eta \nabla \hat{\mathcal{L}}_a(\theta_a^{t-1}) & \hat{\theta}_b^t &= \theta_b^{t-1} - \eta \nabla \hat{\mathcal{L}}_b(\theta_b^{t-1}) \\ \bar{\theta}_a^t &= (1 - \alpha) \cdot \hat{\theta}_a^t + \alpha \cdot \hat{\theta}_b^t & \bar{\theta}_b^t &= (1 - \beta) \cdot \hat{\theta}_b^t + \beta \cdot \hat{\theta}_a^t \\ \theta_a^t &= (1 - I_a^t) \cdot \hat{\theta}_a^t + I_a^t \cdot \bar{\theta}_a^t & \theta_b^t &= (1 - I_b^t) \cdot \hat{\theta}_b^t + I_b^t \cdot \bar{\theta}_b^t \end{aligned}$$

Lemma D.3. Let $X_u^T X_u$ be a p.d. matrix. Let $\lambda_{\max}(X_u^T X_u)$, $\lambda_{\min}(X_u^T X_u)$ be the maximum and minimum eigenvalues of $X_u^T X_u$. Denote the condition number $\rho_u := \frac{\lambda_{\max}(X_u^T X_u)}{\lambda_{\min}(X_u^T X_u)}$, and $Y_u = X_u \theta^*$ then,

$$\frac{1}{\rho_u \Delta_u^t} \hat{\mathcal{L}}_u(\hat{\theta}_u^t) \leq \hat{\mathcal{L}}_u(\bar{\theta}_u^t) \leq \frac{\rho_u}{\Delta_u^t} \hat{\mathcal{L}}_u(\hat{\theta}_u^t) \quad (7)$$

Proof. In the noiseless case, $Y_u = X_u \theta^*$, which gives

$$\hat{\mathcal{L}}_u(\theta) = \|X_u \theta - Y_u\|_2^2 = \|X_u \theta - X_u \theta^*\|_2^2 = \|X_u(\theta - \theta^*)\|_2^2$$

Another way to write this is $\hat{\mathcal{L}}_u(\theta) = (\theta - \theta^*)^T X_u^T X_u (\theta - \theta^*)$. Since $X_u^T X_u$ is a positive definite matrix, we have

$$\lambda_{\min}(X_u^T X_u) \|\theta - \theta^*\|_2^2 \leq \hat{\mathcal{L}}_u(\theta) \leq \lambda_{\max}(X_u^T X_u) \|\theta - \theta^*\|_2^2$$

This implies the following

$$\frac{\hat{\mathcal{L}}_u(\hat{\theta}_u^t)}{\hat{\mathcal{L}}_u(\bar{\theta}_u^t)} \leq \frac{\lambda_{\max}(X_u^T X_u) \|\hat{\theta}_u^t - \theta^*\|_2^2}{\lambda_{\min}(X_u^T X_u) \|\bar{\theta}_u^t - \theta^*\|_2^2} = \rho_u \Delta_u^t$$

and

$$\frac{\hat{\mathcal{L}}_u(\hat{\theta}_u^t)}{\hat{\mathcal{L}}_u(\bar{\theta}_u^t)} \geq \frac{\lambda_{\min}(X_u^T X_u) \|\hat{\theta}_u^t - \theta^*\|_2^2}{\lambda_{\max}(X_u^T X_u) \|\bar{\theta}_u^t - \theta^*\|_2^2} = \frac{\Delta_u^t}{\rho_u}$$

□

Theorem D.4. (*Always Trade, Linear Case*) Let Δ_b^t be the gain-from-trade of agent B at round t . Assume $\Delta_b^t > 1$ for all rounds, making agent B always buy parameters. Let $\delta_b^t = \min_t \Delta_b^t$ and $\delta_b^t / \rho_b > 1$. Then for any $\epsilon \in (0, 1)$ at the end of round T , we have $\hat{\mathcal{L}}_z(\theta_b^T) - \hat{\mathcal{L}}_z(\theta^*) \leq \epsilon$, for

$$T \geq \frac{1}{\log(\delta_b^t / \rho_b)} \log \left(\frac{\hat{\mathcal{L}}_b(\theta_b^0) - \hat{\mathcal{L}}_z(\theta^*)}{\epsilon} \right)$$

Proof. The proof follows by using Lemma D.3 to establish a recurrence relation on $\hat{\mathcal{L}}_b(\theta_b^t)$ and then using the fact that $\hat{\mathcal{L}}_z(\theta) \leq \hat{\mathcal{L}}_b(\theta), \forall \theta$ gives us the result. The steps are as follows,

$$\begin{aligned} \hat{\mathcal{L}}_b(\theta_b^t) &= \hat{\mathcal{L}}_b(\bar{\theta}_b^t) \leq \frac{\rho_b}{\Delta_b^t} \hat{\mathcal{L}}_b(\hat{\theta}_b^t) \\ &\leq \frac{\rho_b}{\Delta_b^t} \left(\hat{\mathcal{L}}_b(\theta_b^{t-1}) - \frac{\eta}{2} \|\nabla \hat{\mathcal{L}}_b(\theta_b^{t-1})\|_2^2 \right) \\ &\leq \frac{\rho_b}{\Delta_b^t} \hat{\mathcal{L}}_b(\theta_b^{t-1}) \end{aligned}$$

The second step follows from the fact that $\hat{\theta}_b^t = \theta_b^{t-1} - \eta \nabla \hat{\mathcal{L}}_b(\theta_b^{t-1})$ and since $\hat{\mathcal{L}}_b$ is L-smooth, this implies the iterates $\{\hat{\theta}_b^t\}$ satisfy descent lemma, which is

$$\hat{\mathcal{L}}_b(\hat{\theta}_b^t) \leq \hat{\mathcal{L}}_b(\theta_b^{t-1}) - \frac{\eta}{2} \|\nabla \hat{\mathcal{L}}_b(\theta_b^{t-1})\|_2^2$$

Now, since $\hat{\mathcal{L}}_z$ is non-negative,

$$\begin{aligned} \hat{\mathcal{L}}_b(\theta_b^t) - \hat{\mathcal{L}}_z(\theta^*) &\leq \frac{\rho_b}{\Delta_b^t} \hat{\mathcal{L}}_b(\theta_b^{t-1}) - \hat{\mathcal{L}}_z(\theta^*) \\ &\leq \frac{\rho_b}{\Delta_b^t} \left(\hat{\mathcal{L}}_b(\theta_b^{t-1}) - \hat{\mathcal{L}}_z(\theta^*) \right) \quad \text{since } \rho_b / \Delta_b^t < 1 \end{aligned}$$

Using the above recurrence we get,

$$\hat{\mathcal{L}}_b(\theta_b^T) - \hat{\mathcal{L}}_z(\theta^*) \leq \left(\prod_{t=1}^T \frac{\rho_b}{\Delta_b^t} \right) \left(\hat{\mathcal{L}}_b(\theta_b^0) - \hat{\mathcal{L}}_z(\theta^*) \right)$$

Since $\delta_b^t = \min_t \Delta_b^t, \forall t$, then for desired error ϵ on the upper bound we get,

$$T \geq \frac{1}{\log(\delta_b^t / \rho_b)} \log \left(\frac{\hat{\mathcal{L}}_b(\theta_b^0) - \hat{\mathcal{L}}_z(\theta^*)}{\epsilon} \right)$$

□

Proposition D.5. (*Never Trade, Linear Case*) Let agents never trade and ignore the market i.e. $I_a^t = 0, I_b^t = 0$ for all t . Suppose they run individual gradient descent on functions $\hat{\mathcal{L}}_a, \hat{\mathcal{L}}_b$ using exact line search. Then for any $\epsilon \in (0, 1)$ at the end of round T , we have $\hat{\mathcal{L}}_z(\theta_b^T) - \hat{\mathcal{L}}_z(\theta^*) \leq \epsilon$, for $T \geq \frac{1}{\log((\rho_b+1)/(\rho_b-1))} \log \left(\frac{\hat{\mathcal{L}}_b(\theta_b^0) - \hat{\mathcal{L}}_z(\theta^*)}{\epsilon} \right)$.

Discussion. First, we note that in both settings (always trade and never trade) the convergence rate is $\mathcal{O}(\log(\frac{1}{\epsilon}))$, which is expected since the loss function getting minimized is strongly-convex. The differences are mainly visible in the constant factors. Most importantly, the leading constant in the result of Theorem D.4 depends on δ'_b — the minimum improvement ratio over T runs. A higher improvement ratio implies a smaller T (i.e. faster convergence) in the worst-case scenario. This justifies the benefit of purchasing the parameters in this setting. In contrast, the leading constant in the result of Proposition D.5 only depends on ρ_b . We further note that $\frac{\rho_b+1}{\rho_b-1} \geq \frac{1}{\rho_b}$, so to claim better rate in Theorem D.4 we would need $\delta'_b \geq \frac{\rho_b^2+\rho_b}{\rho_b-1}$.

E Experimental Details

In this section, we place more details about our experimental setups. All the implementation is open sourced at <https://github.com/SprocketLab/parameter-market>.

E.1 Data Endowment

In the experiments for collaborative setting (Sec. 6.1), agents are limited to collecting a part of dataset on MNIST [33], CIFAR10 [34], and TinyImageNet [35]. For MNIST and CIFAR10, we make agent A collect 10% of data points from class $0 \sim 4$, but obtain full data points from class $5 \sim 9$. For TinyImageNet, agent A collects 10% of data points from class $0 \sim 99$, but obtains full data points from class $100 \sim 199$. Agent B collects data points in the opposite way. We give broker full data points to evaluate agents’ model performance (i.e. $\hat{\mathcal{L}}_z(\theta)$) and to inform gain-from-trade in the try-before-purchase mechanism.

E.2 Model Training Settings

MLP models are trained with 5 hidden layers, each comprising 32 units, using ReLU activations between layers and no normalization. We employ SGD as the optimizer with a learning rate of 10^{-2} . ResNet20 models are trained from scratch and the optimizer used is Adam with a learning rate of 10^{-2} . In both settings, the batch size is set to 256.

F Trading Guidance in the Market

In the following paragraphs, we provide more experimental results. We use the collaborative setting to trade entire parameter sets and apply model alignment. We validate that (i) trading offers instantaneous improvements, (ii) trading after every training epoch results in the fastest convergence, (iii) trading makes agents who obtain fewer data earn more improvements relatively, and (iv) trading asynchronous parameters enhance performance as well. Last but not least, we generalize our setting to a multi-agent market and show parameter trading enables every agent in the market to achieve higher accuracy performance compared to out-of-market training.

How does the epoch to start trading affect the convergence? We study this using two MLPs with MNIST datasets. Figure 6 shows the comparison of agents’ testing loss when they begin trading after different epochs, namely $\{ 10, 20, 30, 40 \}$. Our results indicate that joining the market at *any epoch* provides immediate improvements. Additionally, training parameters and then start trading at an earlier epoch results in faster convergence (i.e. starting to trade after epoch 10 leads to the highest accuracy performance at the end).

How does the trading frequency affect the convergence? Next, we examine how the trading frequency affects convergence. Our experiment asks agents to trade after every $\{ 1, 7, 10 \}$ epoch and compares their testing loss to agents who don’t trade at all. The results, shown in Figure 7, indicate that the optimal approach for agents is to trade after *every epoch*, which leads to the fastest convergence. Despite the fact that testing loss increases during periods when agents don’t trade, they are still able to achieve lower losses overall.

How does data endowment affect trading benefits? We also experiment with different levels of data endowment, ranging from 10% to 80%, and investigate the impact of data endowment on

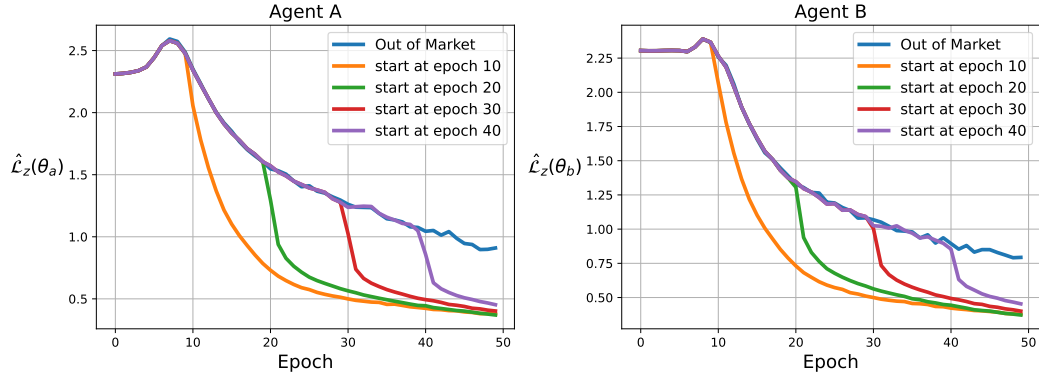


Figure 6: Joining the market at any epoch provides *instantaneous* trading benefits.




Figure 7: Trading parameters after *every epoch* gives the fastest convergence.

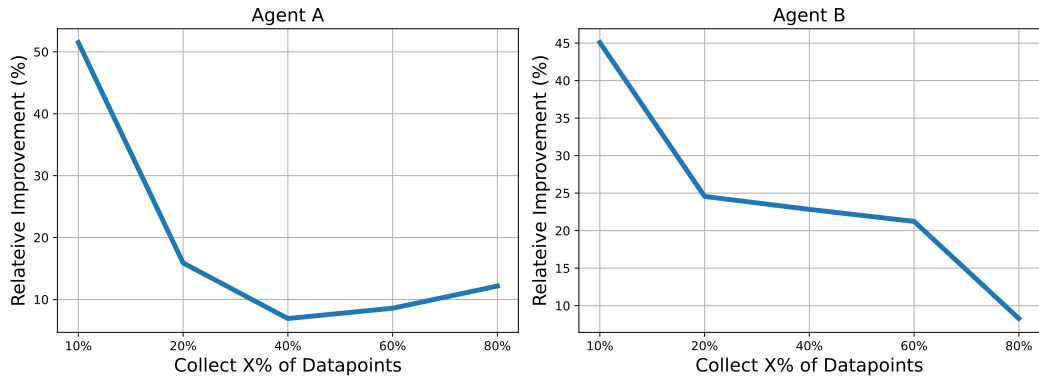


Figure 8: **More limited access to data earns more trading benefits.** Relative improvements are computed by the empirical testing loss compared to out-of-market agents.

trading benefits. Agents are set to collect these percentages of data points for half of the classes. The results, shown in Figure 8, indicate that agents who are more limited in access to data can benefit more from trading relatively. When agent *A* and agent *B* endowed with 10% data points can earn the most trading benefits through improving testing loss by **51.5%** and **45.1%**, respectively.

		Agent A	Agent B
		Testing Acc. (%)	Testing Acc. (%)
MNIST + MLP	out-of-market w alignment	68.50% 88.78%	72.97% 82.51%
CIFAR10 + ResNet20	out-of-market w alignment	71.14% 78.41%	70.56% 73.40%

Table 3: We conduct a new experiment on asynchronous parameter trading. Both agents are asked to train the model for 60 epochs in total. Agent *A* is instructed to delay trading. Results emphasizes the crucial role of brokers in aligning parameters and optimizing purchase weights to eliminate differences not only in synchronous but also in asynchronous parameter trading.

		Agent A	Agent B	Agent C
		Testing Acc. (%)	Testing Acc. (%)	Testing Acc. (%)
MNIST + MLP	out-of-market w alignment	68.07% 82.29%	73.79% 77.08%	76.38% 77.08%
CIFAR10 + ResNet20	out-of-market w alignment	67.62% 79.53%	74.99% 77.77%	74.12% 76.80%
TinyImageNet + ResNet20	out-of-market w alignment	20.51% 32.62%	17.99% 32.07%	18.50% 31.08%

Table 4: We generalize our setting to involve more agents. In the three-agent market, results also follow expectations. The proposed parameter trading is able to help agents achieve higher accuracy performance compared to conventional model training without trading (out-of-market).

How does delayed parameter trading affect the performance? Now we consider another practical scenario that permits asynchronous parameter trading. In this setting, both agents are asked to train the model for 60 epochs in total. Agent *B* trains the model for 5 epochs, trades in the market for 50 epochs, and then trains the model for the remaining 5 epochs. On the other hand, agent *A* is instructed to delay trading. Agent *A* trains the model for 10 epochs and then trades in the market for 50 epochs. In Table 3, our results show that the parameter market allows delayed agent actions to enhance performance as well, with trading in alignment yielding optimal accuracy as expected. This emphasizes the crucial role of brokers in aligning parameters and optimizing purchase weights to eliminate differences *not only in synchronous but also in asynchronous parameter trading*. Note as well that in Table 1, agents train for 5 epochs and then trade synchronously for 55 epochs, which have additional 5 epochs for training then trading. This explains the degraded performance of agent *B*.

How does more agents involved affect trading benefits? For the sake of simplicity and to demonstrate the viability of parameter trading, we present a two-agent market in the main body, which is a commonly used economic model for studying many settings. Nevertheless, it is straightforward to expand the proposed trading framework to include multiple agents. The trading logic can be generalized to look for trades that enable agents to make the largest advancements. This means purchasing parameters that can bring the largest gain-from-trade (Δ_u^t). To validate this, we generalize our setting to involve more agents. We implement a three-agent market by reusing the same data endowment for agent *A* and agent *B* and display performance results in Table 4. We make a third agent, agent *C*, collect 10% of data points from the class 3 ~ 7 in MNIST and CIFAR10, and 10% of data points from class 50 ~ 149 in TinyImageNet. In the three-agent market, results show that the generalization functions as expected: the proposed market is able to help agents achieve higher accuracy performance compared to conventional model training without trading.

G Discussion

In this section, we discuss limitations, (potential) concerns, the needs of broker’s computational infrastructure, and the common prior that we assume in the valuation.

G.1 Limitations

There are two primary limitations toward building viable parameter markets.

First, the reliability of the broker is crucial to the success of the market. The broker must report trade gains and losses without any bias to ensure the market remains efficient. Suppose that the broker acts adversarially, rather than being a neutral third party as intended. In this case, the broker could potentially use models (or their combination) from both parties, could misinform one party or another of their potential benefit and engage in front-running, etc. Hence, the requirement that some entity is willing to serve as an impartial broker is one of the limits of our work.

To achieve this, we can motivate third parties to play this role by earning transaction fee to facilitate trades, or by giving other types of incentives [38]. On the other hand, such a limitation may not be substantial. In fact, having a reliable broker is **essential for any trading market**, just as it is in real-life scenarios (for example, stockbrokers help perform huge numbers of trades, and large-scale escrow organizations help aid in billion-dollar transactions). This similar requirement is accepted and used in practice in other settings, such as (non-machine learning) markets [38].

An additional limitation is that while agents are allowed to train models using various network architectures for a range of downstream tasks, the parameter sets to be traded require a certain alignment. We believe this limitation is not too severe: it can be overcome, for example, by distilling knowledge onto the same space and dimension for merging [39].

G.2 (Potential) Concerns

One potential concern is about model copyright. In fact, one of the motivations for proposing parameter marketplace is precisely *the fact that trading data is more challenging due to the state of digital property rights*. Parameter trading, on the other hand, does not suffer from the same limitations, at least in many present legal frameworks. That is, parameters are typically owned by organizations or individuals training their models. Because of this, voluntarily trading parameters does not require breaking any type of copyright law. Broker is the only additional party to access to these. Besides, we note that our market permits the trading of subsets of parameters (see Sec. 6.1.2) and transferred parameters only if a trade is made, thereby safeguarding a certain degree of model confidentiality.

Another potential concern is how to prevent out-of-bounds parameter usage (i.e., how to ensure a buyer does not publicly release the purchased parameters, sell them to another unauthorized party). There is a wide array of existing research on this question. Applicable tools enable users to check the purchased weights by creating transaction hashes [40] and watermarks [41].

G.3 Broker’s Infrastructure

It is essential to concern the accessibility of computational resources for the broker. In our framework, the broker has two main computational requirements for conducting a trade. Firstly, the broker needs to help agents validate the model’s performance through inference, which is generally less intensive than training [42], and the amount of validation data needed does not need to be high [43]. Second, the broker assists in aligning parameters for both parties using an approximation algorithm for a bilinear assignment problem [44]. This algorithm has been shown to perform efficiently in experiments. Hence, we believe that the computational needs can be met by charging transaction fees to the agents. In other words, the computational workload for the broker is not a significant issue and can be managed.

G.4 Common Priors in Valuation

The common priors when agents are pricing are that buyer values parameters and sets a price arising from gain-from-trade, where we write $v_u(\theta_{u'}^t) = \Delta_u^t$, and this price is derived from a known probability distribution. Both are inspired by Myerson’s auction work [23], designed to help sellers estimate the values of trades. The first prior is reasonable, since all agents consider trading benefits while assessing a parameter trade. The more they stand to gain, the higher the price they are willing to pay. Therefore, to maximize the seller’s revenue, it is best to set the price as close as possible to the buyer’s price. Then in Theorem 4.1, we help agents to bound the other’s valuation. Second, the notion that the buyer’s price is a random variable drawn from a probability distribution is standard and commonly used in economic modeling.