
A Case for Ordinal Peer-evaluation in MOOCs

Nihar B. Shah

U. C. Berkeley

nihar@eecs.berkeley.edu

Joseph K. Bradley

U. C. Berkeley

josephkb@berkeley.edu

Abhay Parekh

U. C. Berkeley

parekh@berkeley.edu

Martin Wainwright

U. C. Berkeley

wainwrig@berkeley.edu

Kannan Ramchandran

U. C. Berkeley

kannanr@berkeley.edu

Abstract

MOOCs have been highly successful due to the ease of disseminating information: anyone with an Internet connection can watch videos of the lectures and download study material. However, they still lag far behind conventional classrooms in one critical aspect—feedback to and evaluation of the students—due to severe mismatches in the number of students enrolled and the number of experts available. One means of performing evaluation and feedback is *peer evaluation* wherein the answers submitted by a student are anonymized and provided to a set of other students to evaluate. As per current peer evaluation techniques, these peer evaluators assign *cardinal* scores to the given solutions. In this paper, we posit an alternative approach to peer evaluation based on pairwise comparisons. We argue that such an *ordinal* approach is significantly more robust to the lack of expertise of the evaluators, as compared to the conventional cardinal approaches.

1 Introduction

The advent of massive open online courses (or MOOCs) via platforms such as Coursera and EdX has enabled millions of people all over the world to gain access to quality education in a cheap and convenient manner. These courses typically have enrollments of a few thousand students, and these numbers are growing at a fast pace. The greatest advantage of MOOCs is the ease of disseminating information: anyone with an Internet connection can watch videos of the lectures and download the study material. However, MOOCs still lag far behind conventional classrooms in one critical aspect: feedback to and evaluation of the students. Due to the massive scale of these courses, it is impossible to have the instructor or the teaching assistants evaluate the thousands of answers. Moreover, since these courses are offered for free (and hopefully will continue to be), hiring paid experts for this task is not a feasible option.

The two most promising alternative evaluation techniques are auto evaluation and peer evaluation, which have been tried with varying degrees of success. Auto evaluation uses software to evaluate students' solutions to homeworks or exams. While auto evaluation techniques are well-suited for questions which are reasonably objective, such as multiple-choice questions or writing computer programs, they do not perform well for subjective questions. Consider, for instance, an essay from a literature class or a smartphone app design from a class on human-computer interaction. These assignments cannot be evaluated well by modern machine learning, necessitating the opinion of a human [1]. Indeed, many fields of study are intrinsically subjective.

An alternative means of performing evaluation and feedback is *peer evaluation*. In a system employing peer evaluation, the answers submitted by a student are anonymized and provided to a set of other students to evaluate. A rubric provided by the instructor serves as a guideline for the students in their evaluation process. For example, Coursera employs peer evaluation for the human-computer



Figure 1: An illustration of the feedback that a student would receive under a scheme of ordinal evaluations. The figure considers a question in which students are required to design a clock app for mobile phones.

interaction (HCI) course [2,3] in which the homework of each student is evaluated by 3 to 5 students, and the final evaluation of a student is computed as the median of these evaluations.

Despite its immense potential, peer evaluation has had only a limited acceptance in MOOCs. Only a handful of courses have attempted to employ peer evaluation. In these courses, the final evaluations obtained by the students under the peer evaluation process were often far from an expert evaluation. Such irregularities have led to a severe opposition to the employment of peer evaluation in MOOCs [4]. The reasons are rather obvious. A significant fraction of students enrolled in the MOOCs are not serious about the course [5], let alone about evaluating their peers' works. The high attrition rates of students also result in complex dynamics in the system. Even when students mean well, lack of expertise and the presence of biases results in significant noise in their evaluations. The system is also faced with additional soft constraints, such as limits on the number of solutions that a student can be asked to evaluate, or existence of a very small pool of expert evaluators who need to be used smartly.

This paper considers the problem of designing peer evaluation schemes. Conventional evaluation schemes typically follow the *cardinal* approach of assigning a *score* to each student (e.g., the HCI course at Coursera). In this paper, however, we argue for the case of an alternative *ordinal* (comparative) approach towards peer evaluation. In such an approach, each student-evaluator will be given some *pairs* of solutions, and in each pair he/she must choose the 'better' solution. The pairs may comprise the solutions provided by other students, or solutions that have been previously calibrated by experts. These comparisons will then be aggregated by a system which assigns a final evaluation to each student. In what follows, we elaborate on this approach, and discuss its possible merits and demerits as compared to the traditional cardinal approach.


The rest of the paper is organized as follows. Section 2 compares the cardinal and ordinal approaches towards peer evaluation. Section 3 discusses our initial work at modeling ordinal peer evaluations and presents some preliminary results. Section 4 provides concluding comments.

2 Cardinal vs. Ordinal Evaluations

The traditional method of evaluation has been cardinal in nature, but such evaluations have been typically performed by experts. It is unclear whether this approach is robust to the lack of expertise. In this section we argue for an alternative ordinal means of evaluation, and provide a qualitative comparison of the two methods.

2.1 Calibration Issues

Every student has inherent biases which may vary with time or depend on the quality of the other answers evaluated. A student may be conservative and always assign moderate scores, while another student may have a tendency to inflate evaluations to the extremes. These inconsistencies are hard to learn with sample sizes as small as 5 or 10 that one would typically get in a peer evaluation setup. Furthermore, these inconsistencies vary with time, thus making them harder to learn and model. On the other hand, such inconsistencies are automatically eliminated (by design) in the ordinal setup. As illustrated below by means of practical experiments, people can be significantly more competent at performing ordinal evaluations than cardinal ones.



Error in Ordinal	12.89%	Error in Ordinal	6.54%
Error in Cardinal	17.45%	Error in Cardinal	13.50%
Additional Ties in Cardinal	0.00%	Additional Ties in Cardinal	15.69%

(a) Estimating age of people from photographs

(b) Estimating areas of circles

Table 1: Screenshots and error-rates of two experiments comparing cardinal and ordinal tasks on Amazon Mechanical Turk. (The screenshots show only the ordinal part of the experiment.)

2.2 Experimental Results on Accuracy in Cardinal vs. Ordinal Evaluations

One could obtain ordinal evaluations from cardinal ones: after collecting cardinal evaluations, any pair of solutions can be compared using the cardinal scores. Such an argument suggests that an ordinal approach does not provide any additional data, and in fact leads to a loss of information. In this section, we present results from some experiments we performed using the Amazon Mechanical Turk crowdsourcing platform that suggest quite the opposite. The experiments reveal that when evaluations are performed by humans, ordinal evaluations contain *significantly less noise* than cardinal evaluations.

(a) *Estimating age of people from photographs*: Each task given to the workers required estimating the age of 10 people whose photographs were shown. In the ordinal setup, the workers had to choose the older person from pairs of pictures, while the cardinal setup required the workers to enter the estimated age. A total of 100 workers performed the tasks.

(b) *Estimating areas of circles*: Each task given to the workers comprised 25 questions. In the cardinal version of the tasks, for each question, the worker was shown a circle in a bounding box, and the worker was required to identify the percentage of the box’s area that the circle occupied. In the ordinal version, the worker was shown two circles in separate, identical bounding boxes, and the worker was required to identify the circle that occupied a larger percentage of area in its respective box. The bounding box was 200 pixels wide and 200 pixels high, and the radius of the circle was chosen as $30 * Beta(15, 3)$, where $Beta$ denotes the Beta distribution. The cardinal answers were then converted to an ordinal form by choosing pairs of questions and looking at which circle was given a higher value. A total of 50 workers performed the tasks.

Table 1 shows the results. Converting cardinal answers to ordinal answers results in a significantly higher error rate than directly asking for ordinal evaluations.

2.3 Ease of Evaluation

The problem of peer evaluation in MOOCs requires dealing with not only the inconsistencies due to lack of expertise, but also with scarcity of resources. It is practically infeasible to ask each student to devote too much of his/her time and effort in the peer evaluation tasks. It is thus of interest to maximize the amount of information that can be gathered under a limited effort by the students, and to design a peer evaluation process that does not fatigue the student-evaluators. To this end, it is fairly well known [6, 7] that humans often find it significantly easier to compare than score. In the setup of peer evaluation, an evaluator may find the task of providing a simple comparison to be much easier than providing a precise numerical score. An ordinal approach would then allow for the collection of more evaluations for the same level of effort, as compared to cardinal evaluations.

2.4 Concrete Constructive Feedback

Given that the evaluators are not experts, we posit that ordinal feedback, as shown in Fig. 1, may often be more desirable to the student as compared to comments on an absolute scale. Since humans are better at comparisons, the evaluator may be able to offer more insightful comments on the positive and negative aspects of the student’s solution, and make these points more concrete by comparing with another solution.¹ Such a pointed evaluation will also guide the student in understanding precisely how he/she can improve his/her solution.

3 Modeling and Inference in Ordinal Peer-evaluation

In this section, we describe preliminary work on modeling and inference in ordinal peer evaluation. After discussing reasonable modeling assumptions, we give a negative result showing that any model matching these assumptions will require non-convex optimization for learning. We then propose a simple model for estimating scores from pairwise comparisons. Despite the non-convexity of the model, initial results on synthetic and real data show promise.

3.1 Modeling and Inference in Ordinal Peer evaluation

We give an axiomatic discussion of models for ordinal peer evaluation. We assume that each of n students $i \in [n]$ has an inherent skill $w_i \in \mathbb{R}$ which we wish to infer using peer comparisons. The peer comparisons let us indirectly observe the skills in two ways: the probability of j beating ℓ depends on their relative skill, and the probability of a correct comparison (where the more skilled student wins) depends on the skill of the peer evaluator.

The peer comparison mechanism works as follows. Pick a set of three students (i, j, ℓ) . Anonymize the answers of students j and ℓ , and ask evaluator i which of the two answers is better. Repeat this process, choosing the students such that every answer is evaluated by some minimum number of students and such that no student performs more than some maximum number of comparisons.²

We can now outline generative models for comparisons. For students i, j, ℓ , we write $(i : j > \ell)$ to denote the event of evaluator i rating j above ℓ . We express a generative model with a function $f : \mathbb{R}^3 \rightarrow [0, 1]$, where

$$f(w_i, w_j, w_\ell) = P(i : j > \ell | w_i, w_j, w_\ell). \quad (1)$$

Note that $P(i : \ell > j | w_i, w_j, w_\ell) = f(w_i, w_\ell, w_j) = 1 - f(w_i, w_j, w_\ell)$. Also, f may depend on other parameters extraneous to our current discussion. From here on, we hide the parameters w_i, w_j, w_ℓ for brevity.

What properties should f have? We state two intuitive axioms:

Axiom 1 Better students are better evaluators. The function $f = P(i : j > \ell)$ must be increasing in w_i if $w_j > w_\ell$ (decreasing if $w_j < w_\ell$).

Axiom 2 Most students i are not malicious. The function $f = P(i : j > \ell)$ must be non-decreasing in w_j and non-increasing in w_ℓ .

Now consider the problem of inferring the skills $\mathbf{w} := [w_1, \dots, w_n]$, given a dataset of peer comparisons, using maximum-likelihood estimation. To permit efficient optimization for inference, we would like our probabilistic model f to be log-concave in \mathbf{w} . Unfortunately, as the following proposition shows, there exists no smooth log-concave function that satisfies our axioms.

Proposition 1 *No twice-differentiable function $f : \mathbb{R}^3 \rightarrow [0, 1]$ satisfies Axioms 1 and 2.*

This result indicates that optimization may always be difficult when modeling ordinal peer evaluation. However, our initial empirical results in the next section are encouraging. In particular, for

¹Note that the other solution chosen for comparison is picked cleverly by the system; the solution may be that provided by a student or may be one that has been calibrated by an expert.

²We assume that the set of triplets $\{(i, j, \ell)\}$ chosen for the comparisons is provided a priori to the algorithm. The problem of evaluator assignment is important, but we omit discussion due to lack of space.

certain models or parameter regions, optimization may be tractable. We are currently working on understanding these questions, with the goal of proving strong guarantees for optimization.

3.2 Refereed Bradley-Terry-Luce (RBTL) Model

We generalize the classical Bradley-Terry-Luce (BTL) model [8, 9] to incorporate the notion of a peer referee. The BTL model says that when two entities j and ℓ with respective skills w_j and w_ℓ are compared, j wins with probability:

$$P_{BTL}(j > \ell) = \frac{1}{1 + \exp(-(w_j - w_\ell))} \quad (\text{BTL model}) \quad (2)$$

We incorporate the idea of peer evaluation by scaling this probability with a function of the evaluator’s skill, giving the Refereed BTL (RBTL) model:

$$P(i : j > \ell) = \frac{1}{1 + \exp(-g_i(w_j - w_\ell))} \quad \text{where } g_i = aw_i + b \quad (\text{RBTL model}) \quad (3)$$

for some parameters a and b . The *evaluation ability* g_i increases linearly with the evaluator’s skill w_i . A skilled evaluator with $g_i \gg 0$ will likely predict $j > \ell$ iff $w_j > w_\ell$; a malicious evaluator with $g_i < 0$ will tend to do the opposite; and a random evaluator with $g_i = 0$ will choose each of j and ℓ with equal probability. Observe that (3) reduces to the original BTL model when $a = 0, b = 1$.

Given a set of comparisons $\{(i : j > \ell)\}$, we can write out the data log-likelihood as a function of \mathbf{w}, a, b . As expected from Proposition 1, this function is not jointly convex in its parameters. We discuss empirical convergence on synthetic data in the next section.

Reasonable parameter regions: Note that the RBTL model only fits Axioms 1 and 2 if $g_i > 0, \forall i$. Axioms 1 and 2 thus constrain our model parameters; if we posit a generative model of the skills w_i , then we can write down constraints for parameters a, b . For example, if $w_i \sim \mathcal{N}(\mu, \sigma^2)$ i.i.d., and $\Phi(\cdot)$ is the standard normal CDF, then we get for each i :

$$P(\text{malicious}) = \Phi\left(-\frac{\mu+b/a}{\sigma}\right). \quad (4)$$

One could imagine many models of ordinal peer evaluation; we make no claim that the RBTL model is ideal. Yet exploring the RBTL model helps us to understand two key questions:

- How hard is optimization for a model with reasonable parameters encoding our axioms?
- How strong is the relation between student skill w_i and evaluator ability g_i ?

3.3 Simulations

To understand the parameterization of the RBTL model, and to test the difficulty of optimization, we ran tests on data simulated from the RBTL model in (3) for various settings of parameters a, b . We generated skills $w_i \sim \mathcal{N}(0, 1)$ and comparisons using the RBTL model. We then trained the RBTL model to estimate the skills \mathbf{w} , fixing a, b to the ground truth. We used stochastic gradient to optimize the data log likelihood, with L2-regularization set according to the prior from which we generated \mathbf{w} . We set the stochastic gradient step size to $1/\sqrt{T}$ on iteration T .

In short, optimization did not pose a significant problem. To give an idea of sample complexity, we have plotted results in Figure 2 for selected models. The error in skill estimates drops quickly as each grader makes more comparisons, but note that many parameter settings still require 10 to 20 comparisons per grader to achieve low error. Interestingly, even though we expect about 8 malicious graders when $a = 1$, estimation is still easy. We posit that large a makes outcomes more certain (i.e., more signal per comparison), even while increasing the number of malicious graders.

3.4 Experiments on Real Data

We tested on data from the third offering of Human Computer Interaction (HCI) on Coursera, taught by Prof. Scott Klemmer (then at Stanford). See [3] for details on the peer evaluation system. In our dataset of peer evaluations for the first English homework assignment, we had 1879 students and 7242 numerical peer evaluations.

We compared our RBTL model with the original BTL model and with median prediction (using the median of a student’s grades as the estimated grade). We tested using 4-fold cross validation on the peer grades, training on 3/4 of the peer grades and testing on the held-out 1/4 of the grades. We




Figure 2: **Skill Estimation Error vs. Number of Comparisons per Student.** Maximum-likelihood estimation for the gBLT model with over 50 students, $b = 1$ and various values of a . A larger value of a implies a stronger dependence between the skill w_i and grading ability g_i of each student; $a = 0$ gives the original BTL model. Error bars show stderr from 5 trials.

	RBTL	BTL	median
mean 0-1 Error	0.241671	0.272198	0.241407
stderr 0-1 Error	0.00538454	0.0107655	0.00509211
mean ProbError	0.285829	0.44362	–
stderr ProbError	0.0058381	0.00130699	–

Figure 3: **Results on HCI Peer evaluations.** The RBTL, BTL, and median model errors for predicting the results of pairwise comparisons of students in held-out peer grades. Discarding the numerical data in peer grades hurts the performance of the BTL model, which does worse than the median prediction. Including grading ability allows the RBTL model to make up for this loss in performance.

created comparison data from peer evaluations by taking each evaluator i and comparing the scores i gave to each pair of peers (j, ℓ) . This produced 16,310 peer comparisons.

We fit the RBTL model by fixing $b = 1$ and estimating w and a by maximizing training data log likelihood, with regularization. We used alternating block coordinate descent, alternating between (a) fixing a and optimizing w with stochastic gradient and (b) fixing w and optimizing a via a line search.

We regularized w with the equivalent of a $\mathcal{N}(1, \sigma^2)$ prior, and we applied L2 regularization to a with parameter λ_a . We chose σ^2 and λ_a via 3-fold cross-validation. We fit the BTL model analogously, except that a was fixed at 0. Cross validation chosen $(\sigma^2, \lambda_a) = (0.5, 100)$ for the RBTL and $\sigma^2 = 100$ for the BTL model.

Training gave an average of $a = 3.38$ for the RBTL model. Recall our discussion of malicious evaluators in (4). Under our prior for w , these settings for σ^2 and (a, b) imply few malicious evaluators: $P(\text{malicious}) \approx 0.0334$.

We evaluated the error by comparing each of our three model’s predictions of pairwise comparisons within the held-out test data. For a given peer comparison $(i : j > \ell)$, our error metric `0-1 Error` has value 0 if the model assigns a higher score to student j than ℓ , and value 1 otherwise. The error metric `ProbError` has value $|P_{\text{model}}(i : j > \ell) - 1|$ for every comparison $(i : j > \ell)$.

We show our initial results in Figure 3. The BTL model does the worst; the RBTL model and the median prediction have similar performance. Modeling grading ability captures an important effect which significantly improves the performance of ordinal models, in terms of both predicting outcomes of pairwise comparisons (`0-1 Error`) and predicting the probability of outcomes (`ProbError`). The large magnitude of $a = 3.38$ selected via cross-validation, relative to $b = 1$, also highlights the value of modeling grading ability.

It is encouraging that the RBTL model can match the performance of a numerical method, even though converting numerical scores to comparisons discards information. The arguments of Section 2 hint at the possibility of obtaining higher-quality ordinal evaluations if the evaluators are directly asked to perform ordinal feedback, as opposed to converting cardinal scores into ordinal evaluations. We posit that, if peer grades were collected as comparisons, our RBTL model would have improved performance.

4 Conclusions

We posit an ordinal approach to peer-evaluation in MOOCs, which we argue is robust to the lack of expertise among the graders. We also present initial work on modeling and analyzing such an ordinal setting, and obtain encouraging results from preliminary experiments on real and synthetic data.

Acknowledgements

The authors would like to thank Jonathan Huang for allowing us to run experiments on the peer-evaluation data from the HCI courses at Coursera.

References

- [1] Professionals against machine scoring of student essays in high-stakes assessment. <http://humanreaders.org/petition/index.php>.
- [2] Chris Piech, Jonathan Huang, Zhenghao Chen, Chuong Do, Andrew Ng, and Daphne Koller. Tuned models of peer assessment in MOOCs. In *International Conference on Educational Data Mining*, 2013.
- [3] C. Kulkarni, K. Pang-Wei, H. Le, D. Chia, K. Papadopoulos, D. Koller, , and S. R. Klemmer. Scaling self and peer assessment to the global design classroom. In *CHI*, 2013.
- [4] Jonathan Rees. Peer grading cant work. <http://www.insidehighered.com/views/2013/03/05/essays-flaws-peer-grading-moocs>.
- [5] Katy Jordan. MOOC completion rates. <http://www.katyjordan.com/MOOCproject.html>.
- [6] William Barnett. The modern theory of consumer behavior: Ordinal or cardinal? *The Quarterly Journal of Austrian Economics*, 6(1):41–65, 2003.
- [7] Neil Stewart, Gordon DA Brown, and Nick Chater. Absolute identification by relative judgment. *Psychological review*, 112(4):881, 2005.
- [8] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, pages 324–345, 1952.
- [9] R Duncan Luce. Individual choice behavior, a theoretical analysis. *Bull. Amer. Math. Soc.* 66 (1960), 259-260 DOI: <http://dx.doi.org/10.1090/S0002-9904-1960-10452-1> PII, pages 0002–9904, 1960.

Appendix

4.1 Proof of Proposition 1

Let us first assume that the function $f(w_i, w_j, w_\ell)$ depends only on the values w_i and $(w_j - w_\ell)$. Let $g = w_i$ and $z = w_j - w_\ell$. Consider the function $s(g, z) := f(w_i, w_j, w_\ell)$. We shall first show that the function s , if twice differentiable, cannot be log-concave.

For $s(g, z)$ to be log-concave, we need

$$-\nabla^2 \log s(g, z) \succeq 0,$$

i.e., we need

$$\frac{1}{s^2} \begin{bmatrix} (\nabla_z s)^2 - s \nabla_z^2 s & \nabla_z s \nabla_g s - s \nabla_{zg}^2 s \\ \nabla_z s \nabla_g s - s \nabla_{zg}^2 s & (\nabla_g s)^2 - s \nabla_g^2 s \end{bmatrix} \succeq 0. \quad (5)$$

A special case of (5) implies that

$$0 \leq [-\nabla_g s \quad \nabla_z s] \frac{1}{s^2} \begin{bmatrix} (\nabla_z s)^2 - s \nabla_z^2 s & \nabla_z s \nabla_g s - s \nabla_{zg}^2 s \\ \nabla_z s \nabla_g s - s \nabla_{zg}^2 s & (\nabla_g s)^2 - s \nabla_g^2 s \end{bmatrix} \begin{bmatrix} -\nabla_g s \\ \nabla_z s \end{bmatrix} \quad (6)$$

$$= \frac{1}{s} (-(\nabla_g s)^2 \nabla_z^2 s - (\nabla_z s)^2 \nabla_g^2 s + 2 \nabla_g s \nabla_z s \nabla_{gz}^2) \quad (7)$$

and hence

$$0 = -(\nabla_g s)^2 \nabla_z^2 s - (\nabla_z s)^2 \nabla_g^2 s + 2 \nabla_g s \nabla_z s \nabla_{gz}^2. \quad (8)$$

Recall that

$$s(g, -z) = 1 - s(g, z) . \quad (9)$$

Noting that (8) must also hold for $-z$, we substitute (9) in (8) to get

$$0 \leq (\nabla_g s)^2 \nabla_z^2 s + (\nabla_z s)^2 \nabla_g^2 s - 2 \nabla_g s \nabla_z s \nabla_{gz}^2 . \quad (10)$$

Along with (8), this implies that

$$0 = (\nabla_g s)^2 \nabla_z^2 s + (\nabla_z s)^2 \nabla_g^2 s - 2 \nabla_g s \nabla_z s \nabla_{gz}^2 . \quad (11)$$

From (5) we also have $(\nabla_z s)^2 \geq s \nabla_z^2 s$ and $(\nabla_g s)^2 \geq s \nabla_g^2 s$. Substituting this in (11) we get

$$s \nabla_{g^2}^2 s \nabla_{z^2}^2 s \leq \nabla_g s \nabla_z s \nabla_{gz}^2 . \quad (12)$$

Using (9) again, we get

$$(1 - s) \nabla_{g^2}^2 s \nabla_{z^2}^2 s \leq - \nabla_g s \nabla_z s \nabla_{gz}^2 . \quad (13)$$

Since $s \in [0, 1]$ we have

$$\nabla_{g^2}^2 s \nabla_{z^2}^2 s \leq 0 . \quad (14)$$

Due to (5) we also need

$$0 \leq \det(\nabla^2 \log s) \quad (15)$$

$$\begin{aligned} &= [(\nabla_z s)^2 (\nabla_g s)^2 - s \nabla_z^2 s (\nabla_g s)^2 - s \nabla_g^2 s (\nabla_z s)^2 + s^2 \nabla_g^2 s \nabla_z^2 s] \\ &\quad - [(\nabla_z s)^2 (\nabla_g s)^2 - 2s \nabla_{gz}^2 s \nabla_g s \nabla_z s + s^2 (\nabla_{gz}^2 s)^2] \end{aligned} \quad (16)$$

$$= s^2 \nabla_g^2 s \nabla_z^2 s - s^2 (\nabla_{gz}^2 s)^2 , \quad (17)$$

where the last expression is obtained due to (11). Consequently,

$$(\nabla_{gz}^2 s)^2 \leq \nabla_g^2 s \nabla_z^2 s . \quad (18)$$

Invoking (14) we get that $\nabla_{gz}^2 s = 0$. It also follows that either $\nabla_g^2 s = 0$ or $\nabla_z^2 s = 0$. Substituting this in (11), we get that at every point (g, z) it must be that either (a) $\nabla_g s = 0$ or (b) $\nabla_z s = 0$ or (c) $\nabla_g^2 s = 0$ and $\nabla_z^2 s = 0$. This contradicts Axioms 1 and 2.

Finally, we return to the general function $f(w_i, w_j, w_\ell)$. If this function is log concave, so is the equivalent function where w_j and w_ℓ are replaced by an invertible linear transformation $(w_j - w_\ell)$ and $(w_j + w_\ell)$. However, the latter is simply the function $s(g, d)$ which we know cannot be log concave. ■