Modeling classroom teaching as optimal communication

Lawrence Liu¹ & Michael C. Frank¹

¹ Stanford University

Author Note

Thanks to NSF BCS #1456077 for support, and to Noah Goodman, Long Ouyang, and Michael Henry Tessler for helpful feedback.

Author contributions: MCF and LL designed models, wrote code, analyzed data, and wrote the paper.

Correspondence concerning this article should be addressed to Michael C. Frank, 450 Serra Mall, Jordan Hall, Stanford, CA 94305. E-mail: mcfrank@stanford.edu

Abstract

A good teacher is a good communicator. But communicating to a large audience can be difficult if audience members have differing preconceptions and hence construe the same message differently. Following this analogy of teaching as communication, we develop a framework for modeling classroom education as optimal communication to a variable audience. We study simple teaching games where teachers provide examples of a target concept to groups of students. Students synthesize these examples with their prior knowledge in order to induce the concept. We consider strategies for managing variability, including ability grouping (“tracking”), reductions in class size, and formative assessment. With known costs on actions, our model can also be extended for decision-theoretic analysis. This model provides a framework for estimating theoretical limits on the utility of educational interventions.

Keywords: classroom teaching, communication, bayesian modeling, class size, formative assessment

Word count: 3540

Modeling classroom teaching as optimal communication

A good teacher is a good communicator, choosing explanations and examples that allow students to learn concepts effectively. But just as a message that is intended for one particular person can be incomprehensible to another, a lesson that is appropriate for one student can be inappropriate for another. The perfect pace for one student may be too fast or too slow for a second student; an analogy that is helpful for one student could be confusing to another. More generally, the broader the audience for a lesson a teacher wants to deliver, the more difficult it can be to ensure that every student learns. In this paper, we explore this intuition formally, creating a framework for understanding the role of student variability in educational outcomes.

Managing student variability is a central issue in educational policy and practice. Should schools group students into classrooms by ability (“tracking”; Glass & Smith, 1979; Slavin, 1989)? Should policymakers invest resources in decreasing class size, allowing for smaller classes with more opportunities to tailor instruction to individual students (Slavin, 1987; Tomlinson, 1999)? When should teachers conduct formative assessments to learn more about the range of knowledge and abilities held by students in a classroom (L. S. Fuchs & Fuchs, 1986; Sadler, 1989)? Each of these questions has been the focus of intense discussion. Randomized trials can provide real-world evidence for the effects of particular interventions on class size or ability grouping (Duflo, Dupas, & Kremer, 2015; e.g., Duflo, Dupas, & Kremer, 2011), but such experiments can tell us only if a particular implementation *does* work, not what could or couldn’t work *in principle*. To our knowledge, no body of theory provides guidance on this issue. Our goal here is to take a first step towards such a framework, which can then be used as a guide for interpreting the assumptions underlying intervention studies.

We exploit the parallel between teaching and communication that emerges in recent formal work using probabilistic models (N. D. Goodman & Frank, 2016; Shafto, Goodman, & Frank, 2012). For example, Shafto, Goodman, and Griffiths (2014) model the process of

reasoning about which examples a teacher should choose to illustrate a concept and what inferences the student should draw from the examples. The unit of their analysis is a “teaching game” (Avrahami et al., 1997; Shafto et al., 2014), by analogy to “signaling games” in linguistics (Lewis, 2008): scenarios whose participants, rules, and parameters are known and hence which admit quantitative analysis. These abstractions can in turn be leveraged to provide insight into more complex, messier situations.

In a teaching game, a teacher attempts to teach a target concept to a group of one or more students by providing examples of that concept. Each student then attempts to induce the target concept based on their prior knowledge and the given examples. Teaching games are collaborative games: the payoff is successful communication such that the student’s learned concept is closer to the target. We formalize this payoff through the information-theoretic notion of information gain, the distance that the student’s concept moves in the direction of the teacher’s through teaching.

The game we study is extremely simple: a teacher attempts to communicate the weight of a biased coin via showing a series of individual coin flip outcomes (heads or tails). The key choice is what set of flips to show. Given unlimited teaching time, a long sequence of flips could allow the learner to approximate the coin’s weight very closely. But with bounded time, the teacher must choose selectively and with respect to the students’ prior beliefs. These priors – which stand in for the combination of knowledge and ability differences that vary across students in real classrooms – are the key factor in determining teaching effectiveness. If two students have very different priors, when they observe the same sequence of examples, they will infer that the teacher is trying to communicate very different weights. Thus, our analysis focuses on managing student variability by intervening on the teaching game.

To explore the limits imposed by structural features of the teaching game, we examine outcomes when both teachers and learners perform optimally. Students in our models are optimal Bayesian learners, who make statistical inferences about how to combine observed examples of a concept with their prior beliefs. Teachers are optimal communicators, who

80 consider their students' prior distribution and choose the best examples to maximize
81 learning. Even under these idealizing assumptions, student learning is necessarily limited.
82 Our analysis investigates sources of these limitations: in particular, the choice of
83 instructional groups, strategies, and teaching examples given uncertainty about student
84 knowledge and abilities (their priors).

85 Uncertainty about student priors plays an important role in our framework. If teachers
86 know exactly what students know – that is, if they know each students' prior beliefs – they
87 can choose an optimal set of examples that will maximize aggregate learning within their
88 classroom. However, in real world classroom settings, it is rare (if not impossible) for
89 teachers to have that exact knowledge of students' prior beliefs. Instead, if teachers are
90 unsure about what students do and do not know, they can guess, informed by observed
91 evidence about student abilities. Such guesses will necessarily be less accurate, leading to
92 lower information gain. We assume that teachers can gather information about student
93 priors by conducting assessments. Such assessments take time away from instruction, but
94 may nonetheless ultimately lead to increased learning due to the information about student
95 priors they provide to teachers. We investigate this tradeoff.

96 We begin by presenting the details of our modeling framework and then describe
97 results that emerge from this framework through simulation. We next describe simulations
98 focused on estimating effects of 1) reducing class size, 2) grouping students into classes by
99 ability, and 3) conducting formative assessments to increase teacher knowledge about
100 students. We end by discussing how our framework can be used for decision-making by
101 assigning costs to individual aspects of an intervention.

102 Though the specifics of quantitative variation in our framework is of course dependent
103 on our choice of teaching game as well as the parameters we choose, a variety of useful
104 generalizations emerge quite clearly. Consideration of how these effects reflect our initial
105 assumptions in turn can provide a framework for interpreting the success and failures of
106 educational interventions.

Model

In our teaching game, students $S = s_1 \dots s_n$ attempt to learn the target concept C , the weight on a biased coin (a single Bernoulli variable). They do so by Bayesian estimation of a sequence of examples $E = e_1 \dots e_m$, using a conjugate Beta-Bernoulli distribution. Their prior distribution is $Beta(\alpha, \beta)$, and the form of the posterior can be written $Beta(\alpha + x, \beta + y)$ where x and y represent the number of heads and tails observed in the data. Under this formulation, the priors α and β represent students' prior knowledge and abilities. If their prior mean $\mu = \frac{\alpha}{\alpha + \beta}$ is distant from C , they will require more examples to estimate C .

Students are sampled so that their beliefs are uniform across the range of possible values for C . Thus, different true values of C represent concepts about which students will have more or less variable priors. For example, for $C = .5$, half of students will have $\mu > C$ and half will have $\mu < C$. A single sample from C (a 0 or 1) will lead half of students to adjust their estimates in the appropriate direction and half to adjust in the wrong direction. For example, showing a 1 (i.e. a heads) will cause all students to increase their posterior means by $\mu_{new} = \frac{\alpha+1}{\alpha+\beta+1}$, but for the half of students with $\mu < C$, they correctly adjust their estimates towards $C = 0.5$, whereas students whose priors $\mu > C$ will incorrectly adjust their estimates away from $C = 0.5$. In contrast, for a more extreme value of C , say $C = .9$, a single sample (likely a 1) will lead students to adjust their estimates up, leading them to be closer to the target concept C which was likely higher than their initial prior means.

In general, the teacher does not know students' priors and must represent them approximately as $\hat{S} = \hat{s}_1 \dots \hat{s}_m$. Teachers use these estimates to compute their objective function, to maximize their estimate of students' *information gain* (IG_T) given a particular set of examples E (a set of heads and tails):

$$IG_T(E) = \sum_{s \in S} D_{KL}(T || \hat{s}) - D_{KL}(T || \hat{s} + E)$$

Intuitively, information gain captures the distance that students travel towards the target concept C due to the examples the teacher chooses. The distance metric is the

Kullback-Leibler divergence ($D_K L$; Cover & Thomas, 2012). Using the fact that both teacher and student beliefs are distributed as Betas, we derive a closed form expression for information gain (see Supplemental Information). Using this objective, they can then choose the examples that maximize their estimate of students' information gain.

We distinguish the teacher's approximation of information gain IG_T that can be computed from the approximations in $\hat{S}$ from the true information gain IG that can be computed via exact knowledge of the students' priors S . In the case where teachers have uncertainty about student priors, they can gain information about students by assessing students, i.e. to update $\hat{S}$ to be closer to S . We model *assessments* as samples from student's true distribution, which teachers then use as Bayesian updates to their internal representations of students. Our general case is also one in which teachers have uncertainty about student priors, but in some simulations we also consider the special case where teachers have perfect knowledge of student beliefs, i.e. $\hat{S} = S$, hence $IG_T = IG$.

Simulations

In our simulations, we investigate how different structural features of the teaching game affect the maximal possible information (the limits on optimal performance). We do this investigation via numerical simulation, choosing a representative set of structural values and exhaustively investigating different choice parameters within these. In particular, we generate random *schools* of 100 students and systematically vary five parameters:

- The number of teachers in the school, from a single class of 100 to 10 classes of 10;
- Out of 12 instructional periods, the number of periods dedicated to formative assessments conducted by each teacher, from 0 (0%) to 12 (100%), inversely trading off against the number of periods dedicated to showing examples E to move student Beta-Bernoulli distributions;
- How extreme the teacher's target concept C is, from a concept for which students have very different starting points relative to the concept (.5) to a more extreme concept for

which students will have more similar starting points (.9; by symmetry this range also covers the range from .5 – .1);

- Whether teachers group students into classes randomly or by ability, with ability grouping implemented as optimally partitioning similar students into classes via teachers’ representations of student knowledge; and
- Whether the teachers have perfect or uncertain knowledge about their students (in the latter case, assessments can be implemented to improve teachers’ estimate of student knowledge).

For each parameter setting, we generate 100 schools, examine the effects of these design choices, and then average performance. All simulations were conducted using the probabilistic programming language **WebPPL** (N. D. Goodman & Stuhlmüller, n.d.). All code and data are available at <http://github.com/mcfrank/teaching>.

Figure 1 shows an example parameter regime (target concept $C = .6$, 5 teachers) and the information gain values that result from different instructional settings within this regime. A number of trends are immediately apparent. First, the overall amount of possible information gain varies as a function of the number of assessments. In the limit (100% of possible instruction time is spent on assessments), students are never taught anything and hence they do not learn at all. Second, there is a substantial boost in total information gain due to ability grouping (red vs. blue) for perfect teachers (dotted lines). Third, and perhaps most interesting, this gain is moderated in uncertain teachers, who must conduct assessments in order to group and to teach students successfully. With no assessments, ability grouping is no better than random grouping. As the number of assessments is increased, however, the groupings approach the optimal (though the total number of instructional opportunities declines, decreasing overall information gain). In this parameter regime, the optimal solution is to group by ability and administer three assessments; this optimal solution is shown by the colored dots.

Figure 1. Average information gain plotted by the number of assessments. Colors show whether students are grouped randomly or by ability, and line-type shows teacher knowledge (uncertain, perfect, or naive). The dots show the number of assessments that produces the maximal information gain.

Information Gain Analyses

In the following analyses, we quantify the consequences of different decisions for student information gain, and give some generalizations that hold across parameter regimes.

Class size reduction yields small effects. We can simulate the number of teachers that are available for a population of students. The number of teachers determines class size, which in turn affects the extent to which teachers can customize instruction to their students. We compare to a baseline in which all students in a school are in the same class. We begin by considering the case of perfect teachers, in order to estimate the maximal possible information gain based on reductions in class size. We compute the proportional

Figure 2. Proportional information gain due reduction in class size, without (left) and with ability grouping (right). Colors show different target concepts. Error bars show 95% confidence intervals computed by bootstrap resampling.

information gain due to class size reduction, compared to the single-class baseline:

$$\frac{IG_{N \text{ teachers}} - IG_{1 \text{ teacher}}}{IG_{1 \text{ teacher}}}.$$

Class size reductions lead to increases in information gain, especially for less extreme concepts, but they are small when they are the only intervention (Figure 2, left). Variance between students in our model is the primary factor governing information gain, with lower variance leading to greater IG. Thus, class size effects in this model are due exclusively to chance reductions in variance. On average, the range of student abilities in a class will be somewhat smaller for a smaller class, leading to an increased ability for the teacher to customize to that class. But this effect is not large, and reaches a maximum of around 10%

of total information gain in our simulations.

Ability grouping into a small number of classes can yield large gains in performance. We next consider the effects of student grouping, in which students are placed in classes based on the teacher’s evaluation of their priors. We compute proportional information gain due to ability grouping similarly to our computation above, except that we compare to the ungrouped condition as a baseline. Figure 2, right shows the results of this analysis. Several trends are apparent. First, all trends are positive. As class sizes decrease, information gain increases over baseline, sometimes substantially (up to 80% in some conditions). Second, gains due to ability grouping are most apparent for less extreme target concepts – that is, the more students differ in their initial knowledge/abilities with respect to these concepts, the more grouping matters. Third, class size plays a relatively small role: beyond two or three teachers, additional gains are marginal. Thus, within our teaching games, our simulations show that – for teachers with perfect knowledge – grouping students into even just a few “tracks” based on their prior knowledge will produce substantial benefits, especially for concepts where students’ starting points are variable.

Imperfect ability grouping yields lower effect sizes, except for highly uncertain concepts. We previously considered teachers with perfect knowledge of students ($\hat{S} = S$). But what if teachers are uncertain about student knowledge and only improve their representation of student knowledge $\hat{S}$ by sacrificing class time to conduct assessments? In this case, each assessment provides a small increment of information to the teacher about the students’ prior distributions (since assessment is modeled as random sampling from students’ true distributions). Necessarily, however, assessment takes away instruction time (as is visible in Figure 1), a tradeoff resulting from a fixed number of instructional periods.

Figure 3 shows proportional information gain due to ability grouping for uncertain teachers. Although perfect ability grouping yields large increases in information gain, these increases are mostly not realized in cases of imperfect knowledge, with one exception. For

Figure 3. Proportional information gain due to grouping using imperfect information, for one target concept (.6). Colors show different numbers of assessments. Error bars show 95% confidence intervals computed by bootstrap resampling.

target concepts above $C = .5$, maximal proportional gains are at most 50%. Only for $C = .5$ are gains larger: this is a case where student priors differ from one another so significantly that it is very difficult to teach without placing students into separate groups. In this case, maximal proportional gains are not achieved until nearly all class time is used for assessment. More modest amounts of assessment lead to much more limited proportional gains.

A small amount of assessment typically produces maximal absolute gains.

How much assessment produces the best possible information gain, optimizing the tradeoff between assessment and instruction? Figure 4 shows the number of assessments that yield the maximal information gain for each target concept, for both for random and ability-grouped class assignments. In both cases, the optimal number of assessments is

Figure 4. Absolute information gain plotted by the optimal number of assessments, for students either grouped by ability or randomly assigned to classes. Colors show different target concepts.

typically relatively limited: in general three or fewer assessments ($< 25\%$ of instructional time) produces optimal results. It is only for $C = .5$, and only for ability grouped classrooms, reinforcing the result above: for target concepts in which student priors are highly variable and a single example can heavily influence a student distribution's accuracy for better or worse, teachers must know a large amount about each individual students' abilities before being able to effectively assign them to appropriate classrooms.

Decision-Theoretic Analyses

In our second set of analyses, we consider an additional factor in making decisions: the costs and payoffs of different interventions. Interventions to change class size are typically

costly: they require resources to be allocated to hire more teachers. By assigning a notional
“cost” C_T (in arbitrary units) to teachers, it is possible to define a Pareto frontier – a set of
maximally efficient parameter settings – between cost and information gain.

Figure 5. Information gain by cost of number of teachers only, with colors showing the number of assessments for each parameter setting. The optimal assessment value for each cost is outlined and connected with a black line to show the Pareto frontier. Subplots show different levels of teacher knowledge and ability grouping. Target concept shown is .6.

Figure 5 shows an example Pareto frontier for one target concept, assuming both
teachers with uncertainty about students and assignment to classrooms by ability. As
described above, for ability grouping, adding a second classroom improves performance
markedly, but returns diminish with additional classrooms. In contrast, when students are
randomly assigned to classes, increasing the number of teachers helps only marginally and a
small number of assessments serves to maximize information gain within each classroom.

Figure 6. Pareto frontiers for three example cost function regimes (panels). Line types show a range of target concepts (only .5, .6, and .7 are shown to decrease the range of values plotted), and point size shows the number of teachers. Other plotting conventions are as in Figure 5.

Other cost structures are possible, however, leading to a number of different possible regimes. Assuming that assessments have a fixed cost C_A , due to classroom-external time or other resources expended in acquiring or creating them, then a number of regimes emerge. Figure 6 shows three sample regimes, in which assessment costs dominate ($C_A > C_T$), teaching costs dominate ($C_T > C_A$), or they are roughly similar ($C_A \approx C_T$), each shown for a range of target concepts. When assessment costs dominate, the most optimal choice is typically to hire many teachers, visually represented by consistently large circles (high teacher count) with varying shades of color (different numbers of assessments) across different target concepts and total cost levels. Interestingly, the converse is not true: even

when teaching costs dominate, a quite small number of assessments is typically optimal. Across all simulations, most gains can be typically realized with a small number of assessments and a small number of teachers due to the diminishing returns of incremental teachers and assessments.¹

General Discussion

What are the limits on classroom learning? We used optimal models of teaching and learning to investigate how structural features of the educational situation restrict learning. In particular, we focused on the role of student variability: The more variable the students in a class are, the less effective any particular lesson is, on average. Our simulations allow us to test the effectiveness of different simulated educational interventions. Within our chosen initial conditions, class size reduction alone yielded only small increases in learning alone. In combination with grouping students into a small number of classes, however, gains increased. These gains were moderated by teachers' knowledge about individual students' priors, though, and were typically optimal when done in combination with a handful of formative assessments to improve the accuracy of teachers' representation of their students' prior knowledge. To assist in decision-making, we added costs to the different actions in our model. Across a range of different cost regimes, most information gain was realized by conducting a small number of assessments and then splitting students by ability into a small number of distinct classes. In sum, this model takes a first step towards creating a framework for investigating theoretical limits on the benefits of educational interventions.

Results from our model could in principle be used to help parse the complex and variable educational literature. The literature on class-size reduction in particular is extremely challenging to navigate, with important studies suggesting null (Hanushek, 1997;

¹ This result mirrors a more general result due to Vul, Goodman, Griffiths, and Tenenbaum (2014): under some sampling cost, a single sample from a binary variable – in our case assessment – is often enough for many decision problems.

Hoxby, 2000), moderate (Duflo et al., 2015; Krueger, 2003), and larger (Mosteller, 1995) effects. Our simulations suggest that there is no single effect of class size reduction that we should expect. Instead, improvements in performance should depend on a number of factors, including initial class size (with reductions more important for classes that are already small, cf. Glass & Smith, 1979) and student variability (with reductions more important for populations that are more variable, cf. Mosteller, 1995) in addition to the complexity of the learning concepts. Rather than resolving conflicting measurements in the literature, our model suggests that different initial conditions should lead interventions to have substantially different effects.

As with any highly abstracted representation of a complex, multifaceted human system, our model has many limitations; we address a small number below. First, our model addresses exclusively “cognitive” aspects of the educational process – the communication of knowledge from teacher to students. It has no provision for the myriad “non-cognitive” factors that are well-known to be extremely important in educational outcomes (Blackwell, Trzesniewski, & Dweck, 2007; e.g., Blair & Razza, 2007; Duckworth, Peterson, Matthews, & Kelly, 2007). One application of models like ours is in understanding the ways that their predictions *differ* from observed data. With further model development, differences between predictions and outcomes could in principle even provide a measure of the relative importance of modeled and unmodeled factors in a particular case.

Second, the model of teaching and learning we assume is direct communication between teacher and student. This model neglects peer effects on instruction (cf. Duflo et al., 2011), active learning by students, and myriad other interactive dynamics that are at play in the classroom. These effects could either dampen or accentuate the effect we investigated. For example, Montessori classrooms have been suggested to be effective in part because they leverage student variability as a strength in student-directed activities (A. S. Lillard, 2016; A. Lillard & Else-Quest, 2006). Modeling peer teaching could be an exciting extension of our framework.

Finally, the teaching game we assume is – by design – exceedingly simple. In order to consider general dynamics of communication in a classroom environment, we have abstracted away completely from the specifics of what is being taught. We view this abstraction as a strength in that it allows for interpretable inferences about the structure of the game without complexities relating to content, but this assumption neglects the important dependencies within and across academic domains. In practice, learning any particular piece of knowledge depends on many others – for example, learning content knowledge often presupposes language and literacy skills – and these dependencies add immense complexity to the implementation of any intervention of the type schematically described here.

Our model instantiates a linkage between a particular set of assumptions and their natural outcomes. The key assumption is that education is fundamentally about the communication of knowledge from a speaker to an audience of variable size and with varying starting points. Once this assumption is adopted, a number of regularities regarding class size, ability grouping, and assessment emerge almost immediately. These linkages suggest future directions for both formal and empirical inquiry – in particular, testing the relationship between quantities approximated in our model and student learning outcomes.

References

- Avrahami, J., Kareev, Y., Bogot, Y., Caspi, R., Dunaevsky, S., & Lerner, S. (1997). Teaching by examples: Implications for the process of category acquisition. *The Quarterly Journal of Experimental Psychology: Section A*, 50(3), 586–606.
- Blackwell, L. S., Trzesniewski, K. H., & Dweck, C. S. (2007). Implicit theories of intelligence predict achievement across an adolescent transition: A longitudinal study and an intervention. *Child Development*, 78(1), 246–263.
- Blair, C., & Razza, R. P. (2007). Relating effortful control, executive function, and false belief understanding to emerging math and literacy ability in kindergarten. *Child*

340 *Development*, 78(2), 647–663.

341 Cover, T. M., & Thomas, J. A. (2012). *Elements of information theory*. John Wiley & Sons.

342 Duckworth, A. L., Peterson, C., Matthews, M. D., & Kelly, D. R. (2007). Grit: Perseverance
343 and passion for long-term goals. *Journal of Personality and Social Psychology*, 92(6), 1087.

344 Duflo, E., Dupas, P., & Kremer, M. (2015). School governance, teacher incentives, and
345 pupil–teacher ratios: Experimental evidence from kenyan primary schools. *Journal of*
346 *Public Economics*, 123, 92–110.

347 Duflo, E., Dupas, P., & Kremer, M. (2011). Peer effects, teacher incentives, and the impact
348 of tracking: Evidence from a randomized evaluation in kenya. *The American Economic*
349 *Review*, 101(5), 1739–1774.

350 Fuchs, L. S., & Fuchs, D. (1986). Effects of systematic formative evaluation: A
351 meta-analysis. *Exceptional Children*.

352 Glass, G. V., & Smith, M. L. (1979). Meta-analysis of research on class size and
353 achievement. *Educational Evaluation and Policy Analysis*, 1(1), 2–16.

354 Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic
355 inference. *Trends in Cognitive Sciences*, 20(11), 818–829.

356 Goodman, N. D., & Stuhlmüller, A. (n.d.). *The design and implementation of probabilistic*
357 *programming languages*. Retrieved 2017, from <http://dippl.org>

358 Hanushek, E. A. (1997). Assessing the effects of school resources on student performance:
359 An update. *Educational Evaluation and Policy Analysis*, 19(2), 141–164.

360 Hoxby, C. M. (2000). The effects of class size on student achievement: New evidence from
361 population variation. *The Quarterly Journal of Economics*, 115(4), 1239–1285.

362 Krueger, A. B. (2003). Economic considerations and class size. *The Economic Journal*,
363 113(485).

364 Lewis, D. (2008). *Convention: A philosophical study*. John Wiley & Sons.

365 Lillard, A. S. (2016). *Montessori: The science behind the genius*. Oxford University Press.

366 Lillard, A., & Else-Quest, N. (2006). The early years: Evaluating montessori education.

367 *Science*, 313(5795), 1893–1894.

368 Mosteller, F. (1995). The tennessee study of class size in the early school grades. *The Future*
369 *of Children*, 113–127.

370 Sadler, D. R. (1989). Formative assessment and the design of instructional systems.

371 *Instructional Science*, 18(2), 119–144.

372 Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: The
373 consequences of psychological reasoning for human learning. *Perspectives on Psychological*
374 *Science*, 7(4), 341–351.

375 Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical
376 reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55–89.

377 Slavin, R. E. (1987). Ability grouping and student achievement in elementary schools: A
378 best-evidence synthesis. *Review of Educational Research*, 57(3), 293–336.

379 Slavin, R. E. (1989). Class size and student achievement: Small effects of small classes.
380 *Educational Psychologist*, 24(1), 99–110.

381 Tomlinson, C. A. (1999). *The differentiated classroom: Responding to the needs of all*
382 *learners*. Association for Supervision & Curriculum Development.

383 Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal
384 decisions from very few samples. *Cognitive Science*, 38(4), 599–637.