> *Therefore it seemeth to me, that the truest way to understand conversation is to know the faults and errors to which it is subject, and from thence every man* [sic] *to form maxims to himself* [sic] *whereby it may be regulated.*
> —Jonathan Swift (1713), Hints Towards an Essay on Conversation

Introduction

Teachers have the potential to greatly influence the learning of their students; this much is indisputable. And, as we look for ways to improve the quality of education for students in the United States, it is inevitable that our attention will turn to teachers and the ways they influence student learning. Most of us can recall examples from our own schooling of "good teachers" in whose classrooms we learned a great deal. That is, we have a notion that there are more and less effective ways to practice the craft of teaching. Our efforts to measure teacher effectiveness reflect the pursuit of a noble, simple, and decidedly democratic goal: every student should be taught by an effective teacher.

This goal is easy to state, but its pursuit has led to an enormous, complex, impassioned national conversation about how to measure teacher effectiveness. It is not the purpose of this paper to eliminate the fundamental differences of belief and opinion that undergird this conversation, nor to mitigate the passion with which these beliefs and opinions are held, nor to shed new light on some technical or arcane aspect of the debate. Instead, the purpose of this text is to invite readers to step outside of the conversation momentarily. By examining the conversation about teacher effectiveness from an unfamiliar angle, we can better understand which language behaviors within this conversation are functional and which are counterproductive. A conceptual framework developed by media theorist and cultural critic Neil Postman (1976), *communication as a semantic environment*, provides a useful tool for doing just this.

When analyzing particular language behaviors, it is useful to consider the context in which they occur. First, the three contextualizing components of the semantic environment will be considered: *people, purposes,* and *rules of discourse*. Then, specific language behaviors that lead to conflict in the semantic environment of "measuring teacher effectiveness" will be analyzed. The hope is that analyzing these behaviors will allow us to reduce instances of disagreement to those occasions when there truly is a substantive and important point on which to disagree, avoiding those occasions when language unsuited to its context has obfuscated the real issue at hand.

The Semantic Environment

Communication, conceptualized broadly, includes a wide range of information exchanges, from a conversation between colleagues, to research published in an academic journal, to a political advertisement, to federal legislation, to a newspaper editorial, to a blog post, to a boisterous happy hour debate.

Postman (1976) proposed that communication is best understood and analyzed as an event that occurs within a *semantic environment*, which is characterized by three interrelated parts: people, their purposes, and the general rules of discourse by which such purposes are achieved. These three features of the semantic environment surround and influence the actual language behaviors that occur in any communicative situation; this relationship is illustrated in Figure 1.

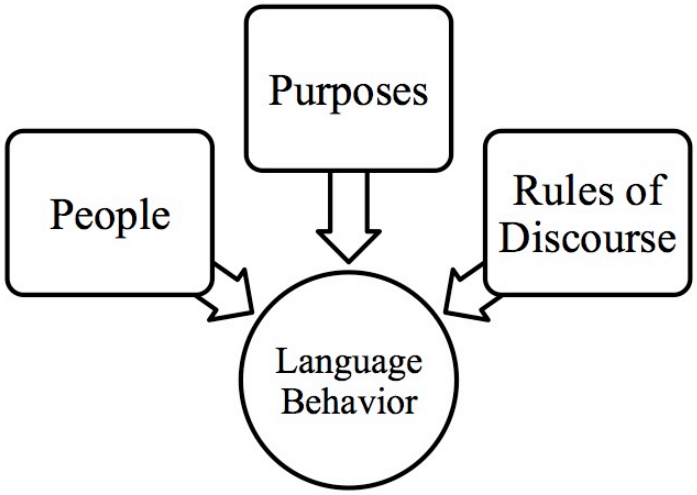

Figure 1. This diagram depicts the three features of the semantic environment (people, purposes, and rules of discourse) and their relationship to actual language behaviors.

Semantic environments can be thought of as social structures in which people interact in order to do something to, for, with, or against other people. Semantic environments can be nested: a principal observing and scoring a teacher's lesson is a semantic sub-environment nested within the larger semantic environment of school-based teacher evaluation. They can evolve over time (e.g. the semantic environment of the Elementary and Secondary Education Act). They can be wide in their influence (e.g., the semantic environment of student-centered pedagogy, in which discussions range from how to help students reflect on their work to the utility of the Reading Workshop approach to language arts

instruction) or entirely local and short-lived (e.g., the semantic environment of a particular parent-teacher conference, in which the focus is solely on the progress of one student). Currently in the United States, the broad semantic environment of teacher evaluation is rife with confusion and conflict, and many of the examples used in this paper are drawn from that environment.

It makes little sense to try to analyze the value or reasonableness of a language behavior without knowledge of the context in which that behavior occurred. Or, put differently, a germane and useful statement in one semantic environment could be useless or harmful in another. The statement, "National Board Certification is only very weakly correlated with student achievement data," functions differently at a research symposium on measuring teacher practice than it does at a party in celebration of a co-worker's receipt of such certification. Because the utility of a remark is determined by the totality of the situation in which it is made, the three contextualizing components of the semantic environment of measuring teacher effectiveness—people, purposes, and rules of discourse—will be considered in turn.

People

Insofar as questions of measuring teacher effectiveness are questions about the quality of education afforded to the next generation, these questions should concern everyone in our society. Those stakeholders with the more prominent roles in the semantic environment, however, include

- teachers, parents, and students;
- school and district administrators;
- state and federal lawmakers;
- state and federal education agencies and departments;
- research funding agencies;
- publishers of educational curricula and assessments;
- advocates of various educational policies and reforms;
- union members and leaders; and
- those associated with teacher education programs (traditional as well as alternative certification).

Purposes

The purposes for which people participate in a particular semantic environment of measuring teacher effectiveness are as wide-ranging as the types of stakeholders listed above. Parents might attend a school board meeting to advocate for the job security of their child's favorite teacher. A school principal

might attend that same meeting to explain how the teacher evaluation process works at his school. A district official might be there to introduce forthcoming district policies around the use of VAM (value-added modeling) in teacher evaluations. And finally, a political candidate might put in an appearance to glad-hand, network, and prove (or photograph, at least) her interest in "local issues."

Postman (1976) notes that "one of the principal reasons why people are forever quarreling about the quality and relevance of their remarks is that semantic environments are multipurposed" (p. 21), and until we understand the purposes of the participants in a given environment, we are in no position to evaluate the utility of their language choices. As we saw above, even in a semantic environment as localized as a school board meeting, there can be a confluence of several purposes, resulting in a situation susceptible to confusion, conflict, and miscommunication: our principal may consider the parents' remarks too anecdotal and subjective to be relevant; the parents might see the district official's presentation as an excuse to fire their child's teacher; and our politician, finding the whole debate a bit too animated for a good photo-op, might attempt to mitigate disagreement through conciliatory language, which everyone else will see as deliberate muddying of the waters.

When considering the different purposes that motivate different interlocutors, assuming that one's goal is to facilitate productive discourse, the most useful thing is neither to judge the merit of different purposes nor even to define them explicitly. The most useful thing is to recognize that they exist and to understand that they fundamentally influence the way people use (and evaluate the use of) language. Bearing this in mind is an important first step in avoiding counterproductive language behaviors.

A principal tension in the semantic environment of measuring teacher effectiveness arises from two importantly different ways to use the information collected, that is, two different purposes behind the data collection:

- "To identify areas of improvement for individual teachers, leading to the preparation of individual improvement plans (including professional development)" (OECD, 2009, p. 19). These are *formative* purposes, which give teacher evaluation an improvement function.
- "To determine career advancement, award performance rewards, or establish sanctions for underperforming teachers" (OECD, 2009, p. 19). These are *summative* purposes, which give teacher evaluation an accountability function.

Several researchers have noted that there are difficulties in designing a single teacher evaluation system that combines both the improvement function and the accountability function (OECD, 2009; OECD, 2013; Klinger, Shulha, & DeLuca,

2008). While some have concluded that these purposes are too divergent to coexist effectively and have argued for just one of them—"the emphasis of teacher assessment policies should be to make teachers better, not to fire them" (IES, 2012, p. 13)—many see the possibility of "an overall system that ties together evaluation and improvement" (IES, 2012, p. 11). Regardless of which argument one finds more compelling, this much is clear: there are divergent stakeholder beliefs about the relative importance of these two purposes and the feasibility of realizing them with a single system of teacher evaluation. This divergence is a defining feature of the *purpose* component of the semantic environment of measuring teacher effectiveness.

Rules of Discourse

Each semantic environment, Postman (1976) argues, is subject to rules of discourse. Nearly always implicit, these rules govern the vocabulary, types of statements, and methods of claim substantiation that are considered permissible within a given environment. For illustration, consider religious discourse and scientific discourse; suggested descriptions of the rules governing each of these follow in Figure 2.

Religious Discourse	Scientific Discourse
• Religious discourse in the public arena is the use of a certain rhetorical style, a style that conforms to certain rules of underlying structure. This structure includes definition through metaphor; intentional redundancy and ambiguous meanings; repeated drawing of identification between readers and characters in the text; decisive, rather than tentative, statements; and extensive provision of exemplary stories. (Wuthnow, 1988)	• Since the main purpose of the semantic environment called science is to produce reliable and predictable knowledge about the world, the rules of scientific discourse are fairly precise. We are obliged, for example, to put our statements in such a form that they are either verifiable or refutable. We are also obliged to have our statements open to public scrutiny, to express ourselves at all times tentatively, to define our terms concretely, and to keep our language relatively free of ambiguity. (Postman, 1976)

Figure 2. Although usually tacit and frequently debated, rules of discourse shape every semantic environment. Brief summaries of the rules of religious and scientific discourse are presented here.

Source: Postman, 1976; Wunthnow 1988.

Similar summaries could be generated for rules that govern political or legal discourse, or the discourse of advertising, advocacy, or patriotism. No set of rules is inherently better than another; none leads inevitably to a more estimable type of communication. Religious language, as described in Figure 2, can appeal to our emotions and aesthetic sensibilities in a way that scientific language cannot. On the other hand, if our goal is to express our ideas with precision and clarity, scientific rather than religious language may be more useful. Different types of discourse are suited to different purposes, with different merits, and different liabilities. The important thing is that language behaviors that are permissible and efficacious under one set of rules may not be so under another.

Questions about appropriate discourse exist in many semantic environments within the social sciences, where inquiry typically concerns both academics and nonacademics, researchers and practitioners, the highly informed, and the general public. Should the rules of discourse for conversations about measuring teacher effectiveness more closely resemble the rules for political discourse or scientific discourse? Is there a place for advocacy? Advertising?

It is beyond the scope of this paper to establish which types of discourse are more or less appropriate in the semantic environment of measuring teacher effectiveness. In fact, this broad semantic environment has many sub-environments, which are and ought to be governed by different rules of discourse. For example, conversations about the test-retest reliability coefficient of a particular classroom observation protocol are most fruitfully conducted using the type of scientific discourse described above. In contrast, conversations about the potential effects of privately funded research agendas seem to call for political or historical discourse. In short, there are many productive types of discourse within this broad semantic environment. Conflict arises, however, when interlocutors using different sets of rules meet in a single semantic sub-environment. For an example of this conflict, see Figure 3**,** which contains an account of two education researchers who took issue with the type of discourse used by the Secretary of Education in a report on teacher quality.

A Case of Contested Rules of Discourse

In July, 2002, the U.S. Secretary of Education's Annual Report on Teacher Quality (U.S. Department of Education) was released. In December of that year, Educational Researcher published a critique of this report (Darling-Hammond & Youngs, 2002), which criticized the report on many fronts: inaccurate synthesis of existing research, misleading assertions, cherry-picking of evidence, and citation of dubious research.

One wonders if, in addition to these grave accusations, the authors of the critique also objected to the vocabulary used in the report, possibly finding it too emotionally evocative to feature in a report from the Department of Education.

In the leading paragraphs of the critique, the authors quote the report four times, and all of these quotations contain highly charged, emotionally resonant vocabulary:

> Stating that current teacher certification systems are '*broken*,' and that they impose '*burdensome requirements*' for education coursework that make up 'the bulk of current teacher certification *regimes*,' the report argues that certification should be redefined… making student teaching and attendance at schools of education optional and eliminating 'other *bureaucratic hurdles.'* [emphasis added] (Darling-Hammond & Youngs, 2002, p. 13)

While the critique never explicitly references the vocabulary of the report, it continues to quote this type of emotional language throughout, perhaps indicating disagreement about the rules of discourse that ought to govern reports from the Secretary of Education.

Figure 3. This example illustrates a case of contested rules of discourse; the disagreement is between education researchers and the U.S. Secretary of Education.

Source: Darling-Hammond & Youngs, 2002.

Language Behaviors

Having considered the three contextualizing features of the semantic environment of measuring teacher effectiveness, we can turn now to specific language behaviors in this environment. Postman (1976) developed a taxonomy of seventeen specific language behaviors that can function to impede, rather than stimulate, productive discourse. Nearly all of these behaviors are present in contemporary discourse about measuring teacher effectiveness. In order to show how language choices shape perception, policy, and practice, four of the seventeen language behaviors that are particularly prevalent and problematic are discussed below.

Definition Tyranny

Postman (1976) notes that miscommunication can arise from differences in definitions. He describes "people who have so internalized a definition that they cannot even imagine an alternative way of seeing matters; they make *a* definition into *the* definition" (Postman, 1976, p. 188). Figure 4 contains a discussion of how definition tyranny has appropriated, narrowed, and ossified our general understanding of the terms *student achievement, student outcomes, achievement gains,* and *student learning*. Definition tyranny is also evident in the recent, heated exchange between Randi Weingarten, president of the American Federation of Teachers and the advocacy group, Chiefs for Change, a coalition of former and current state education superintendents.

A Case of Definition Tyranny

Within the semantic environment of measuring teacher effectiveness, a particularly pervasive form of definition tyranny exists regarding the terms *student achievement*, *student outcomes*, *achievement gains*, and *student learning*. To an outsider, these may appear to be very nebulous constructs: *What is 'student learning'? How could you hope to quantify it?*

There is consensus in the field, however, that these are to be understood, in roughest terms, as "some sort of desirable performance standardized achievement tests."

Some researchers have gone a step further and defined *teacher effectiveness* as *the ability to bring about some sort of desirable performance by one's students on standardized achievement tests*. For an exemplar of this type of definition tyranny, see The Education Trust's report on the use of VAM in Los Angeles Unified School District (Education Trust West, 2012).

Figure 4. This example shows how terms that may be ambiguously understood in general usage can be very narrowly defined within a particular semantic environment.

Source: Education Trust West, 2012.

Weingarten argued that it is unfair to attach high-stakes consequences (in particular, consequences for teacher evaluations) to new standardized assessments aligned to the Common Core State Standards before teachers have had time to properly absorb and create curriculum around the standards. She called for a moratorium, not on the assessments themselves, but on the use of their results in high stakes decisions (Rose, 2013).

The Chiefs for Change responded with an open letter to Education Secretary Arne Duncan, in which they interpreted Weingarten's recommendation in light of a single, unstated but unwavering, definition of *accountability*: "consequences and rewards based on student progress" (Chiefs for Change, 2013). Excerpts from the letter include:

- "Recently, some members of the national education community have advocated for pulling back on *accountability* [emphasis added] in our schools" (para. 4)
- "The members of Chiefs for Change reject any calls for a moratorium on *accountability* [emphasis added]" (para. 5)
- "A one-size-fits-all suspension of *accountability* [emphasis added]" measures…" (para. 5)

There is legitimate discussion, and even disagreement, to be had about how to ensure that teachers are doing right by students. Nevertheless, this example points out that in the semantic environment of measuring teacher effectiveness, there is little consensus on the definition of *accountability*. This makes sense because the term describes an abstract concept (like *truth* or *fairness*), and so people use the term in different ways. Weingarten herself declines to define *accountability* explicitly; she simply suggests (in her own open letter) that it is part of a much longer process than the one conceived of by the Chiefs for Change, a process that involves standards, curriculum, classrooms, feedback, and improvement (Weingarten, 2013).

I do not claim that simply by recognizing the definition tyranny present in this exchange, we can resolve the differences of opinion that underlie the disagreement. We can, however, change the nature of the conversation. Instead of exchanging open letters in which different implicit definitions function to malign opposing perspectives, we can focus the conversation more closely on the fundamental issue: how, as we introduce a new generation of standards and assessments, can we ensure that teachers are helping their students learn?

Model Muddles

Postman (1976) argues that every semantic environment is controlled by metaphors through which people interpret the meaning and value of statements and actions in the environment. The metaphors to which he refers are deeply ingrained and they shape our thinking and language. For example, a teacher's relationship to school quality may be conceptualized thusly: "The role of an individual teacher in a school is like a player on a football team: all teachers are vital, but the culture of the school is even more important for the quality of the school" (Sahlberg, 2013, para. 10). This conceptualization will influence how one approaches such questions as *should we seek to explain differences in student outcomes in terms of teacher quality or school quality?* and *do student/ parent surveys or the VAM scores of individual teachers provide more valuable information about a school?*

The metaphors directing our thinking and language remain unexpressed in most types of discourse. When people in the semantic environment of measuring

teacher effectiveness operate with different metaphors for teaching, for example, confusion or conflict can arise. Delandshere and Petrosky (1998) discuss two different (and they argue, contradictory) approaches to recording classroom observations: numerical ratings and interpretive summaries. The conflict between these two approaches, they assert, arises from different underlying metaphors of teaching. Assigning a rating involves "recognizing an event as belonging to a category of events" (Delandshere and Petrosky, 1998, p. 21). The unseen metaphor behind this approach is that teaching is a collection of events, each of which aligns (or fails to align) with a sort of "ideal event." In contrast, "the interpretive summary approach seems to be more compatible with the philosophy of teaching […] that recognizes teaching as constructed in context and content, given particular students, and within different learning communities" (Delandshere and Petrosky, 1998, p. 22).

It is, of course, rare for people to analyze and delineate so thoughtfully a source of philosophical conflict within the semantic environment of measuring teacher effectiveness. Often, misunderstanding and disagreement arise because the metaphors influencing a person's thinking and language remain unstated. Any attempt to eliminate conflicting metaphors is probably doomed to fail: people will always conceptualize human relationships and processes differently. But by being explicit about our metaphors and avoiding the assumption that others necessarily share them, we can better locate sources of confusion or conflict.

Propaganda

Postman (1976) devotes several pages to disambiguating the term *propaganda*, which has such contentious connotations and varied meanings. The definition he selects, which is useful for the present analysis, is "language that invites us to respond emotionally, emphatically, more or less immediately, and in an either-or manner . . . distinct from language which stimulates curiosity, reveals its assumptions, causes us to ask questions, invites us to seek further information and to search for error" (Postman, 1976, p. 170). In short, propaganda is language that says "believe this" instead of "consider this." Propaganda often assumes that a question is closed and that all that is required is collective action; this assumption is exemplified in a magazine advertisement from ExxonMobile, depicted in Figure 5.

A Case of Propaganda

ExxonMobil has joined the education reform discussion, publishing a glossy, full page ad in the company's quarterly publication, *The Lamp*. The ad is promoting the National Math and Science Initiative's UTeach program, an Alternative Certification Program (ACP) supported by ExxonMobil. The text appears in a blue sky above a sun that shines out in rays of yellow pencils and beakers, evoking science classrooms:

> **Let's unlock the brilliance of 10,000 skilled teachers**.
> By 2020, the National Math and Science Initiative's Uteach program will have helped more than 10,000 undergrads earn both a degree in math or science and a teaching certificate. They'll do it without spending extra time or money. Those highly skilled teachers will reach an estimated 4 million students nationwide. Join ExxonMobil in supporting programs like UTeach that raise the bar in math and science. Let's invest in our teachers so they can inspire our students. Let's solve this^{SM} (ExxonMobil, 2012).

The growing popularity of ACPs for teacher training and certification is understandable: they provide a broad and relatively inexpensive labor force and, for many who believe teacher quality is the key policy lever to pull in education reform, they seem to point to a feasible and powerful solution. Their expanded presence, however, has brought these programs to the attention of those who study measures of teacher effectiveness. Different ACPs measure and certify their teachers differently, and there is little consensus among researchers, scholars, practitioners, and policymakers about the advisability of expanding such programs (see Darling-Hammond & Youngs, 2002 for a discussion of the key disagreements in this debate).

ExxonMobile clearly appreciates the rhetorical power of the ad's tagline, "Let's solve this^{SM}": they have secured proprietary rights to this phrase as a service mark (SM) in their advertising campaign. An associated website, letssolvethis.com, contains several more examples of the sort of propaganda illustrated above: simple, powerful language that implies that no further study of ACPs is needed, but only the collective will to act.

Figure 5. ExxonMobil has joined the education reform discussion; their use of propaganda in this semantic environment is illustrated here.

Source:ExxonMobil 2012.

There are semantic environments in which the rules of discourse permit, even encourage, propaganda (the semantic environment of commercial advertising, for example). That is, propaganda is not, in and of itself, a problem. It becomes a problem when it masquerades as something else or when it is used in semantic environments whose rules of discourse exclude it.

In 2010, the *Los Angeles Times* began a series entitled "Grading the Teachers," whose central focus was the use of VAM (value-added models) in evaluating teachers. The articles are written in the style of a thoughtful, careful analysis of the broad issue of measuring teacher effectiveness: perspectives of various stakeholders are presented, statistical findings are discussed, and leading researchers are quoted. That is, the series is fashioned along the lines of scholarly research. Within several of the articles, however, we find statements of a type not typically permitted under the rules of discourse associated with scholarly research, statements that typify propaganda (as defined above):

- "Value-added analysis offers the closest thing available to an objective assessment of teachers" (Felch, Song, & Smith, 2013, para. 23). Instead of inviting the reader to seek further information on what is still very much an open question in the semantic environment of measuring teacher effectiveness (i.e., What measurement tools provide the most objective assessment of teachers?), the authors proceed as if the question is closed.
- In reference to a call for the use of value-added scores as one measure of performance, the head of the State Board of Education is quoted as saying, "I think it's simply a failure of will" (Felch, Song, & Smith, 2013, Jobs with Security section, para. 20), and this perspective permeates the series as a whole. Consensus within the semantic environment of measuring teacher effectiveness on the use of value-added scores in teacher evaluation systems does not exist. There is widespread disagreement among leading scholars and researchers on many aspects (technical, conceptual, political, pedagogical, social) of this issue. Far from a closed question that simply requires the will to act, the issue continues to be one of the most widely discussed and researched topics in the field.

Silent Questions

Two people in a semantic environment may believe they are addressing the same question when, in fact, they are addressing different questions. Silent questions, Postman (1976) explains, can pose a threat to effective communication because they hide behind what appears to be consensus. Consider, for example, the large question: which measurement tools are the most accurate indicators of

Silent Questions: a Short Quiz

"More can be done to improve education by improving the effectiveness of teachers than by any other single [within school] factor" (Wright, Horn, & Sanders, 1997, p. 63)."

Research findings such as this have inspired many policy makers to ask, "How can we improve the effectiveness of teachers?" Two different silent questions lurk behind this primary question, and Josh and Joel have each responded to one of them.

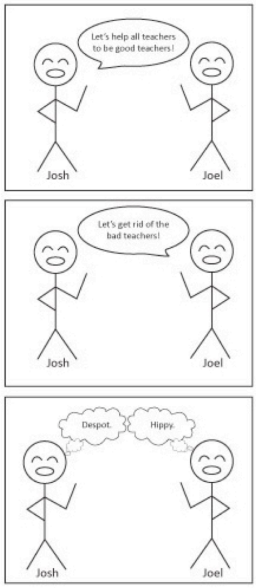

(1) Which silent question has Josh heard and responded to?

(2) Which silent question has Joel heard and responded to?

Figure 6. Different silent questions can lead to different policy recommendations. This phenomenon is illustrated here, around the issue of improving teacher effectiveness.

Source: Wright, Horn, & Sanders, 1997.

teacher effectiveness? Divergent answers to this question arise, in part, from different silent questions lurking behind the large question:

- Which measurement tools can best arrange teachers on a normal distribution?
- Which measurement tools provide the most personalized, locally relevant information?
- Which measurement tools best capture "effective teaching"?

Note that all of these silent questions are worth asking and there is legitimate disagreement on their answers. Imagine, however, a conversation between two people who are unaware that they are answering different silent questions from the list above. Both speakers will evaluate the other's claims and evidence relative to his or her own silent question, and this misalignment will doom effective communication, supplanting it with confusion and doubt about the other's capacity to reason logically.

Figure 6 includes an additional example of misaligned silent questions and suggests that not only can this misalignment derail effective communication, it can also undermine one's respect for a conversation partner. Recognizing silent questions will not eliminate fundamental disagreement on substantive issues, but it will allow people to fruitfully disagree instead of talking past each other as they answer different questions.

Conclusion

The preceding analysis of the semantic environment of measuring teacher effectiveness (and specific language behaviors that confound this environment) is an exercise in what Postman (1976) refers to as *meta-semantics*. The fundamental goal of meta-semantics is to put

oneself outside the context of a given semantic environment, in order to see it in its entirety. With this perspective, it is possible to assess the nature and quality of language behaviors with a relatively high degree of detachment because one is concerned less with participating in the conversation and more with observing it. Postman notes that this shift from "a participant to a participant-observer position is almost always accompanied by a lessening of fervor, a suspicion of ideology, a willing suspension of belief, and a heightening of interest in the process of communication" (p. 237). Insofar as this paper has facilitated, at least momentarily, such a shift in perspective, it has achieved its purpose.

References

Chiefs for Change. (2013, May 21). An open letter of support for Common Core. Retrieved from http://chiefsforchange.org/chiefs-for-change-an-open-letter-of-support-for-common-core/

Darling-Hammond, L. & Youngs, P. (2002). Defining "highly qualified teachers": What does "scientifically-based research" actually tell us? *Educational Researcher, 31*(9), 13-25.

Delandshere, G. & Petrosky, A. R. (1998). Assessment of complex performances: Limitations of key measurement assumptions. *Educational Researcher, 27*(2), 14-24.

Education Trust West (2012, January). Learning denied: The case for equitable access to effective teaching in California's largest school district. Retrieved from http://www.edtrust.org/sites/edtrust.org/files/ETW%20Learning%20Denied%20Report_0.pdf

ExxonMobile. (2012). The Lamp—2012, #2. Retrieved from http://cdn.exxonmobil.com/~/media/The%20Lamp/2012/news_pub_lamp_2012-2.pdf

Felch, J., Song, J. & Smith, D. (2010, August 14) Grading the teachers: Who's teaching L.A.'s kids? *The Los Angeles Times*. Retrieved from http://www.latimes.com/news/local/la-me-teachers-value-20100815,0,258862,full.story

IES. (2012, August 9). Proceedings from IES Meeting: *Learning from recent advances in measuring teacher effectiveness.* Retrieved from http://ies.ed.gov/director/pdf/measuring_teacher_effectiveness.pdf

Klinger, D. A., Shulha, L. M., & DeLuca, C. (2008). Teacher evaluation, accountability, and professional learning: The Canadian perspective. *Pensamiento Educativo, 43*, 209-222.

OECD. (2009, December). *Teacher evaluation: A conceptual framework and examples of country practices.* Retrieved from http://www.oecd.org/education/school/44568106.pdf

OECD. (2013). *Teachers for the 21st century: Using evaluation to improve teaching.* Retrieved from http://www.oecd.org/site/eduistp13/TS2013%20Background%20Report.pdf

Postman, N. (1976). *Crazy talk, stupid talk: How we defeat ourselves by the way we talk—and what to do about it.* New York, NY: Delacorte Press.

Rose, M. (2013, May 1). AFT calls for moratorium on Common Core consequences. Retrieved from http://www.aft.org/newspubs/news/2013/043013commoncore.cfm

Sahlberg, P. (2013, May 15). What if Finland's great teachers taught in U.S. schools? *The Washington Post.* Retrieved from http://www.washingtonpost.com/blogs/answer-sheet/wp/2013/05/15/what-if-finlands-great-teachers-taught-in-u-s-schools-not-what-you-think/

Swift, J. (1713). *Hints towards an essay on conversation*. Retrieved from http://www.ucc.ie/celt/published/E700001-016/index.html

U.S. Department of Education. (2002). *Meeting the highly qualified teachers challenge: The Secretary's annual report on teacher quality.* Washington, DC: U.S. Department of Education, Office of Postsecondary Education, Office of Policy, Planning, and Innovation.

Weingarten, R. (2013, May 21). An open letter to Chiefs for Change. American Federation of Teachers. Retrieved from http://www.aft.org/pdfs/press/chiefsforchange052113.pdf

Wright, S.P., Horn, S.P., & Sanders, W.L. (1997). Teacher and classroom context effects on student achievement: Implications for teacher evaluation. *Journal of Personnel Evaluation in Education, 1*(1), 57-67.

Wuthnow, R. (1988). Religious discourse as public rhetoric. *Communication Research, 15*(3), 318-338.